arXiv:1511.08724v2 [cs.LG] 5 Jan 2016

On the convergence of cycle detection for navigational reinforcement learning

Tom J. Ameloot* Jan Van den Bussche

Abstract

We consider a reinforcement learning framework where agents have to navigate from start states to goal states. We prove convergence of a cycle-detection learning algorithm on a class of tasks that we call reducible. Reducible tasks have an acyclic solution. We also syntactically characterize the form of the final policy. This characterization can be used to precisely detect the convergence point in a simulation. Our result demonstrates that even simple algorithms can be successful in learning a large class of nontrivial tasks. In addition, our framework is elementary in the sense that we only use basic concepts to formally prove convergence.

1 Introduction

Reinforcement Learning (RL) is the subfield of Artificial Intelligence concerned with agents that have to learn a task-solving policy by exploring state-action pairs and observing rewards (Sutton and Barto, 1998). Off-policy algorithms such as Q-learning, or on-policy algorithms such as Sarsa, are well-understood and can be shown to converge towards optimal policies under quite general assumptions. These algorithms do so by updating, for every state-action pair (s, a), an estimate $Q(s, a)$ of the expected value of doing a in s.

Our aim in this article is expressly *not* to propose a more efficient or more powerful new RL algorithm. In contrast, we want to show that convergence can occur already with very simplistic algorithms. The setting of our result is that of tasks where the agent has to reach a goal state in which a reward action can be performed. Actions can be nondeterministic. We like to refer to this setting as *navigational learning*.

The learning algorithm we consider is for a simplistic agent that can only remember the states it has already visited. The algorithm is on-policy; its only update rule is that, when a state is revisited, the policy is revised and updated with an arbitrary new action for that state. We refer to this algorithm as the *cycle-detection algorithm*. Our main result is that this algorithm converges for all tasks that we call *reducible*. Intuitively, a task is reducible if there exists a policy that is guaranteed to lead to reward. We also provide a test for convergence that an outside observer could apply to decide when convergence has happened, which can be used to detect convergence in a simulation. We note that the final policy is allowed to explore only a strict subset of the entire state space.

*T.J. Ameloot is a Postdoctoral Fellow of the Research Foundation – Flanders (FWO).

A first motivation for this work is to understand how biological organisms can be successful in learning navigational tasks. For example, animals can learn to navigate from their nest to foraging areas and back again (Geva-Sagiv et al., 2015). Reward could be related to finding food or returning home. As in standard RL, the learning process might initially exhibit exploration, after which eventually a policy is found that leads the animal more reliably to reward. In the context of biologically plausible learning, Frémaux et al. (2013) make the following interesting observations. First, navigational learning is not restricted to physical worlds, but can also be applied to more abstract state spaces. Second, the formed policy strongly depends on the experiences of the agent, and therefore the policy is not necessarily optimal. We elaborate these observations in our formal framework. We consider a general definition of tasks, which can be used to represent both physically-inspired tasks and more abstract tasks. Furthermore, we do not insist on finding (optimal) policies that generate the shortest path to reward, but we are satisfied with learning policies that avoid cycles.

A secondary motivation for this work is to contribute towards filling an apparent gap that exists between the field of Reinforcement Learning and the more logic-based fields of AI and computer science. Indeed, on the structural level, the notion of task as used in RL is very similar to the notion of *interpretation* in description logics (Baader et al., 2010), or the notion of *transition system* used in verification (Baier and Katoen, 2008). Yet, the methods used in RL to establish convergence are largely based on techniques from numerical mathematics and the theory of optimization. Our aim was to give proofs of convergence that are more elementary and are more in the discrete-mathematics style common in the above logic-based fields, as well as in traditional correctness proofs of algorithms (Cormen et al., 2009).

Standard RL convergence proofs assume the condition that state-action pairs are visited (and thus updated) infinitely often, see e.g. (Watkins and Dayan, 1992). Conditions of this kind are known as *fairness* conditions in the theory of concurrent processes (Francez, 1986). Also for our convergence proof we need an appropriate fairness assumption to the effect that when the agent repeats some policy-updating configuration infinitely often, it must also explore all possible updates infinitely often.

We note that the cycle-detection learning algorithm could be remotely related to biologically plausible mechanisms. In some models of biological learning (Potjans et al., 2011; Frémaux et al., 2013), a policy is represented by synaptic connections from neurons encoding (the perception of) states to neurons encoding actions. Connections are strengthened when pre-before-post synaptic activity is combined with reward (Schultz, 2013), causing an organism to remember action preferences for encountered states. If an organism would initially have a policy that frequently leads to cycles in the task, there is a (slow) way to still unlearn that policy, as follows.[1] Consider a pair (q, a) of a state q and its preferred action a in the policy. Due to noise (Maass, 2014), a neuron x participating in the encoding of action a could become activated just before state q effectively occurs. Possibly, this post-before-pre synaptic activity leads to long-term-depression (Gerstner et al., 2014), i.e., connections from q to x are

[1] In this discussion, we purposely do not mention mechanisms of disappointment, i.e., the opposite of reward, because the framework in this article does not contain such mechanisms.

weakened.[2] At some synapses, the weakening effect is aided by a longer time window for long-term-depression compared to long-term-potentiation (Markram et al., 2011). So, if reward would remain absent for longer periods, as in cycles without reward, noise could gradually unlearn action preferences for states. In absence of such preferences, noise could generate random actions for states. The unlearning phase followed by new random action proposals, would resemble our cycle-detection algorithm.

Outline This article is organized as follows. We discuss related work in Section 2. We formalize important concepts in Section 3. We present and prove our results in Section 4. We discuss examples and simulations in Section 5, and we conclude in Section 6.

2 Related Work

Some previous work on reinforcement learning algorithms is focused on learning a policy efficiently, say, using a polynomial number of steps in terms of certain input parameters of the task (Kearns and Singh, 2002; Brafman and Tennenholtz, 2002; Strehl et al., 2009). There is also a line of work in reinforcement learning that is not necessarily aimed towards efficiently bounding the learning time. In that case, convergence of the learning process happens in the limit, by visiting task states infinitely often. Some notable examples are temporal-difference learning (Sutton, 1988; Sutton and Barto, 1998) and Q-learning (Watkins, 1989; Watkins and Dayan, 1992). Temporal-difference learning has become an attractive foundation for biological learning models (Potjans et al., 2011; Frémaux et al., 2013; Schultz, 2013, 2015).

Most previous works in numerical reinforcement learning try to find optimal policies, and their related optimal value functions (Sutton, 1988; Watkins and Dayan, 1992; Dayan, 1992; Dayan and Sejnowski, 1994; Jaakkola et al., 1994; Tsitsiklis, 1994); an optimal policy gives the highest reward in the long run. This has motivated the design of numerical learning techniques. The corresponding proof techniques do not always clearly illuminate how properties of the task state space interplay with a particular learning algorithm. With the framework introduced in this article, we hope to shed more light on properties of the task state space, in particular on the way that paths could be formed in the graph structure of the task. Although our graph-oriented framework has a different viewpoint compared to standard numerical reinforcement learning, we believe that our Theorem 4.1, showing that convergence always occurs on reducible tasks, and its proof contribute to making the fascinating idea of reinforcement learning more easily accessible to a wider audience. Our convergence result has a similar intent as previous results showing that numerical learning algorithms converge with probability one.

Various papers study models of reinforcement learning in the context of neuronal agents that learn to navigate a physical environment (Vasilaki et al., 2009; Potjans et al., 2011; Frémaux et al., 2013). Interestingly, Frémaux et al.

[2] If neuron x is activated by noise just before state q occurs, refractoriness could prevent state q from subsequently activating x (Gerstner et al., 2014). The resulting absence of a postsynaptic spike at x fails to elicit long-term-potentiation, i.e., connections from q to x are not strengthened. So, the mentioned weakening effect is not compensated.

(2013) study both physical and more abstract state spaces. As an example of a physical state space, they consider a navigation task in which a simulated mouse has to swim to a hidden platform where it can rest, where resting corresponds to reward; each state contains only the x and y coordinate. As an example of an abstract state space, they consider an acrobatic swinging task where reward is given when the tip of a double pendulum reaches a certain height; this space is abstract because each state contains two angles and two angular velocities, i.e., there are four dimensions. Conceptually it does not matter how many dimensions a state space has, because the agent is always just seeking paths in the graph structure of the task.

This idea of finding paths in the task state space is also explored by Bonet and Geffner (2006), in a framework based on depth-first search. Their framework has a more global perspective where learning operations have access to multiple states simultaneously and where the overall search is strongly embedded in a recursive algorithm with backtracking. Our algorithm acts from the local perspective of a single agent, where only one state can be observed at any time.

As remarked by Sutton and Barto (1998, p. 104), a repeated theme in reinforcement learning is to update the policy (and value estimation) while the agent visits states. This theme is also strongly present in the current article, because for each visited state the policy always remembers the lastly tried action for that state. The final aim for convergence, as studied in this article, is to eventually not choose any new actions anymore for the encountered states.

The notion of reducibility discussed in this article is related to the principles of (numerical) dynamic programming, upon which a large part of reinforcement learning literature is based (Sutton and Barto, 1998). Indeed, in reducibility, we defer the responsibility of obtaining reward from a given state to one of the successor states under a chosen action. This resembles the way in dynamic programming that reward prediction values for a given state can be estimated by looking at the reward prediction values of the successor states. In settings of standard numerical reinforcement learning, dynamic programming finds an optimal policy in a time that is worst-case polynomial in the number of states and actions. This time complexity is also applicable to our iterative reducibility procedure given in Section 3.1.

3 Navigational Reinforcement Learning

We formalize tasks and the notion of reducibility in Section 3.1. Next, in Section 3.2, we use an operational semantics to formalize the interaction between a task and our cycle-detection learning algorithm. In Section 3.3, we define convergence as the eventual stability of the policy. Lastly, in Section 3.4, we impose certain fairness restrictions on the operational semantics.

3.1 Tasks and Reducibility

Tasks To formalize tasks, we use nondeterministic transition systems where some transitions are labeled as being immediately rewarding, where reward is only an on-off flag. Formally, a *task* is a five-tuple

$$T = (Q, Q_0, A, \mathit{rewards}, \delta)$$

where Q, Q_0, and A are nonempty finite sets; $Q_0 \subseteq Q$; *rewards* is a nonempty subset of $Q \times A$; and, δ is a function that maps each $(q, a) \in Q \times A$ to a nonempty subset of Q. The elements of Q, Q_0, and A are called respectively *states*, *start states*, and *actions*. The set *rewards* tells us which pairs of states and actions give immediate reward. Function δ describes the possible successor states of applying actions to states.

Remark 3.1. Our formalization of tasks keeps only the graph structure of models previously studied in reinforcement learning; essentially, compared to finite Markov decision processes (Sutton and Barto, 1998), we omit transition probabilities and we simplify the numerical reward signals to boolean flags. We do not yet study negative feedback signals, so performed actions give either reward or no reward, i.e., the feedback is either positive or neutral. In our framework, the agent can observe states in an exact manner, which is a commonly used assumption (Sutton and Barto, 1998; Kearns and Singh, 2002; Brafman and Tennenholtz, 2002; Bonet and Geffner, 2006). We mention negative feedback signals and partial information as topics for further work in Section 6. □

Reducibility Let T be a task as above. We define the set

$$\mathit{goals}(T) = \{q \in Q \mid \exists a \in A \text{ with } (q, a) \in \mathit{rewards}\}.$$

We refer to the elements of $\mathit{goals}(T)$ as *goal states*. Intuitively, for a goal state there is an action that reliably gives immediate reward. Each task has at least one goal state because the set *rewards* is always nonempty. The agent could learn a strategy to reduce all encountered states to goal states, and then perform a rewarding action at goal states. This intuition is formalized next.

Let $V \subseteq Q$. We formalize how states can be reduced to V. Let $\mathbb{N}_0$ denote the set of natural numbers without zero. First, we define the infinite sequence

$$L_1, L_2, \ldots$$

of sets where $L_1 = V$, and for each $i \geq 2$,

$$L_i = L_{i-1} \cup \{q \in Q \mid \exists a \in A \text{ with } \delta(q, a) \subseteq L_{i-1}\}.$$

We call L_1, L_2, etc, the *(reducibility) layers*. We define $\mathit{reduce}(T, V) = \bigcup_{i \in \mathbb{N}_0} L_i$. Note that $\mathit{reduce}(T, V) \subseteq Q$. Because Q is finite, there is an index $n \in \mathbb{N}_0$ for which $L_n = L_{n+1}$, i.e., L_n is a *fixpoint*. Letting $V' \subseteq Q$, we say that V' is *reducible to* V if $V' \subseteq \mathit{reduce}(T, V)$. Intuitively, each state in V' can choose an action to come closer to V. We also say that a single state $q \in Q$ is *reducible to* V if $q \in \mathit{reduce}(T, V)$.

Now, we say that task T is *reducible (to reward)* if the state set Q is reducible to $\mathit{goals}(T)$. We use the abbreviation $\mathit{reduce}(T) = \mathit{reduce}(T, V)$ where $V = \mathit{goals}(T)$. Reducibility formalizes a sense of solvability of tasks.

We illustrate the notion of reducibility with the following example.

Example 3.2. We consider the task $T = (Q, Q_0, A, \mathit{rewards}, \delta)$ defined as follows: $Q = \{1, 2, 3\}$; $Q_0 = \{1\}$; $A = \{a, b\}$; $\mathit{rewards} = \{(3, a), (3, b)\}$; and,

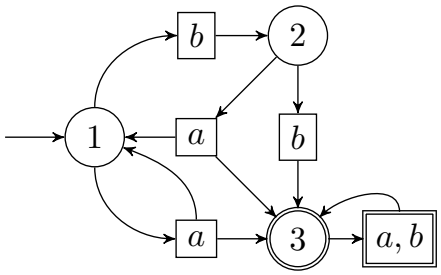

Figure 3.1: The task from Example 3.2. States and action applications are represented by circles and boxes respectively. Start states are indicated by an arrow without origin. Goal states and their rewarding actions are highlighted by a double circle and double box, respectively.

regarding δ, we define

$$\begin{aligned} \delta(1,a) &= \{1,3\}, \\ \delta(1,b) &= \{2\}, \\ \delta(2,a) &= \{1,3\}, \\ \delta(2,b) &= \{3\}, \\ \delta(3,a) = \delta(3,b) &= \{3\}. \end{aligned}$$

Task T is visualized in Figure 3.1. Note that the task is reducible, by assigning the action b to both state 1 and state 2. The reducibility layers up to and including the fixpoint, are:

$$\begin{aligned} L_1 &= \mathit{goals}(T) = \{3\}, \\ L_2 &= \{3,2\}, \\ L_3 &= \{3,2,1\}. \end{aligned}$$

For simplicity, the assignments $1 \mapsto b$ and $2 \mapsto b$ form a deterministic strategy to reward. But we could easily extend task T to a task in which the strategy to reward is always subjected to nondeterminism, by adding a new state 4 with the new mappings $\delta(1,b) = \{2,4\}$, $\delta(4,a) = \delta(4,b) = \{3\}$. □

Remark 3.3. Reducibility formalizes the intuition of an acyclic solution. This appears to be a natural notion of solvability, even in state graphs that contain cycles (Bonet and Geffner, 2006).

We would like to emphasize that reducibility is a notion of progress in the task transition graph, but it is not the same as determinism because each action application, i.e., transition, remains inherently nondeterministic. We may think of reducibility as onion layers in the state space: the core of the onion consists of the goal states, where immediate reward may be obtained, and, for states in outer layers there is an action that leads one step down to an inner layer, closer to reward. When traveling from an outer layer to an inner layer, the nondeterminism manifests itself as unpredictability on the exact state that is reached in the inner layer. □

3.2 Cycle-detection Algorithm

We describe a cycle-detection learning algorithm that operates on tasks, by means of an operational semantics that describes the steps taken over time. We first give the intuition behind the cycle-detection algorithm, and then we proceed with the formal semantics.

3.2.1 Intuition

We want to formally elaborate the intuition of path learning. Our aim therefore is not necessarily to design another efficient learning algorithm. It seems informative to seek only the bare ingredients necessary for navigational learning. How would such a simple algorithm look like?

As a first candidate, let us consider the algorithm that is given some random initial policy and that always follows the policy during execution. There would not be any exploration, and no learning, since the policy is always followed and never modified. In general, the policy might not even lead to any reward at all, and the agent might run around in cycles without obtaining reward.

At the opposite end of the spectrum, there could be a completely random process, that upon each visit to a task state always chooses some random action. If the agent is lucky then the random movement through the state space might occasionally, but unreliably, lead to reward. There is no sign of learning here either, because there is no storage of previously gained knowledge about where reward can be obtained.

Now we consider the following in-between strategy: the algorithm could only choose random actions when it detects a cycle in the state space before reaching reward. If the agent does not escape from the cycle then it might keep running around indefinitely without ever reaching reward. More concretely, we could consider a cycle-detection algorithm, constituted by the following directives:

- Starting from a given start state, we continuously remember all encountered states. Each time when reward is obtained, we again forget about which states we have seen.
- Whenever we encounter a state q that we have already seen before, we perform some random action, and we store that action in the policy (for that state q).

The cycle-detection algorithm is arguably amongst the simplest learning algorithms that one could conceive. The algorithm might be able to gradually refine the policy to avoid cycles, causing the agent to eventually follow an acceptable path to reward. The working memory is a set containing all states that are visited before obtaining reward. The working memory is reset whenever reward is obtained.

3.2.2 Operational Semantics

We now formalize the cycle-detection algorithm. In the following, let $T = (Q, Q_0, A, \mathit{rewards}, \delta)$ be a task.

Configurations A *configuration* of T is a triple $c = (q, \pi, w)$, where $q \in Q$; π maps each $q \in Q$ to an element of A; and $w \subseteq Q$. The function π is called the *policy*. The set w is called the *working memory* and it contains the states that are already visited during the execution, but we will reset w to $\emptyset$ whenever reward is obtained. We refer to q as the *current state* in the configuration, and we also say that c *contains* the state q. Note that there are only a finite number of possible configurations. The aim of the learning algorithm is to refine the policy during trials, as we formalize below.

Transitions We formalize how to go from one configuration to another, to represent the steps of the running algorithm. Let $c = (q, \pi, w)$ be a configuration. We say that c is *branching* if $q \in w$; this means that configuration c represents a revisit to state q, and that we want to generate a new action for the current state q. Next, we define the set $opt(c)$ as follows: letting $A' = A$ if c is branching and $A' = \{\pi(q)\}$ otherwise, we define

$$opt(c) = \{(a, q') \mid a \in A' \text{ and } q' \in \delta(q, a)\}.$$

Intuitively, $opt(c)$ contains the *options* of actions and successor states that may be chosen directly after c. If c is branching then all actions may be chosen, and otherwise we must restrict attention to the action stored in the policy for the current state. Note that the successor state depends on the chosen action.

Next, for a configuration $c_1 = (q_1, \pi_1, w_1)$ and a pair $(a, q') \in opt(c_1)$, we define the *successor configuration* $c_2 = (q_2, \pi_2, w_2)$ that results from the application of (a, q') to c_1, as follows:

- $q_2 = q'$;
- $\pi_2(q_1) = a$ and $\pi_2(r) = \pi_1(r)$ for all $r \in Q \setminus \{q_1\}$; and,
- $w_2 = w_1 \cup \{q_1\}$.

We emphasize that only the action and visited-status of the state q_1 is modified, where q_1 is the state that is departed from. We denote the successor configuration as $apply(c_1, a, q')$.

A *transition* t is a four-tuple (c_1, a, q', c_2), also denoted as $c_1 \xrightarrow{a, q'} c_2$, where $c_1 = (q_1, \pi_1, w_1)$ and $c_2 = (q_2, \pi_2, w_2)$ are two configurations, $(a, q') \in opt(c_1)$, and $c_2 = apply(c_1, a, q')$. We refer to c_1 and c_2 as the *source configuration* and *target configuration*, respectively. We say that t is a *reward transition* if $(q_1, a) \in rewards$. Note that there are only a finite number of possible transitions because there are only a finite number of possible configurations.

Trials and Runs A *chain* $\mathcal{C}$ is a sequence of transitions where for each pair (t_1, t_2) of subsequent transitions, the target configuration of t_1 is the source configuration of t_2. Chains could be finite or infinite.

A *trial* is a chain $\mathcal{C}$ where either *(i)* the chain is infinite and contains no reward transitions; or, *(ii)* the chain is finite, ends with a reward transition, and contains no other reward transitions. To rephrase, if a trial is finite then it ends at the first occurrence of reward; and, if there is no reward transition than the trial must be infinite.

In a trial, we say that an occurrence of a configuration is *terminal* if that occurrence is the last configuration of the trial, i.e., the occurrence is the target configuration of the only reward transition. Note that an infinite trial contains no terminal configurations.

A *start configuration* is any configuration $c = (q, \pi, w)$ where $q \in Q_0$ and $w = \emptyset$; no constraints are imposed on the policy π.

Now, a *run* $\mathcal{R}$ on the task T is a sequence of trials, where

1. the run is either an infinite sequence of finite trials, or the run consists of a finite prefix of finite trials followed by one infinite trial;
2. the first configuration of the first trial is a start configuration;
3. whenever one (finite) trial ends with a configuration $c_1 = (q_1, \pi_1, w_1)$ and the next trial starts with a configuration $c_2 = (q_2, \pi_2, w_2)$, we have *(i)* $q_2 \in Q_0$; *(ii)* $\pi_2 = \pi_1$; and, *(iii)* $w_2 = \emptyset$;[3] and,
4. if the run contains infinitely many trials then each state $q \in Q_0$ is used at the beginning of infinitely many trials.

We put condition (3) in words: when one trial ends, we start the next trial with a start state, we reuse the policy, and we again reset the working memory. By resetting the working memory, we forget which states were visited before obtaining the reward. The policy is the essential product of a trial. Condition (4), saying that each start state is used at the beginning of infinitely many trials, expresses that we want to learn the whole task, with all possible start states.

To refer to a precise occurrence of a trial in a run, we use the ordinal of that occurrence, which is a nonzero natural number.

Remark 3.4. In the above operational semantics, the agent repeatedly navigates from start states to goal states. After obtaining immediate reward at a goal state, the agent's location is always reset to a start state. One may call such a framework episodic (Sutton and Barto, 1998). We note that our framework can also be used to study more continuing operational processes, that do not always enforce a strong reset mechanism from goal states back to remote start states. Indeed, a task could define the set of start states simply as the set of all states. In that case, there are runs possible where some trials start at the last state reached by the previous trial, as if the agent is trying to obtain a sequence of rewards; but we still reset the working memory each time when we begin a new trial. □

3.3 Convergence

We now define a convergence property to formalize when learning has stopped in a run. Consider a task $T = (Q, Q_0, A, \mathit{rewards}, \delta)$. Let $\mathcal{R}$ be a run on T.

Definition 3.5. We say that a state $q \in Q$ *(eventually) becomes stable* in $\mathcal{R}$ if there are only finitely many non-terminal occurrences of branching configurations containing q.

[3] Note that c_2 satisfies the definition of start configuration.

An equivalent definition is to say that after a while there are no more branching configurations at non-terminal positions containing q. Intuitively, eventual stability of q means that after a while there is no risk anymore that q is paired with new actions, so q will definitely stay connected to the same action.[4] Note that states appearing only a finite number of times in $\mathcal{R}$ always become stable under this definition.

We say that the run $\mathcal{R}$ *converges* if *(i)* all trials terminate (with reward), and *(ii)* eventually all states become stable. We say that the task T is *learnable* if all runs on T converge.

Remark 3.6. In a run that converges, note that the policy will eventually become fixed because the only way to change the policy is through branching configurations at non-terminal positions. The lastly formed policy in a run is called the final policy, which is studied in more detail in Section 4.2. We emphasize that a converging run never stops, because runs are defined as being infinite; the final policy remains in use indefinitely, but it is not updated anymore.

We would also like to emphasize that in a converging run, eventually, the trials contain no cycles before reaching reward: the only moment in a trial where a state could be revisited, is in the terminal configuration, i.e., in the target configuration of the reward transition. □

3.4 Fairness

There are two choice points in each transition of the operational semantics:

- if the source configuration of the transition is branching, i.e., the current state is revisited, then we choose a new random action for the current state; and,
- whenever we apply an action a to a state q, we can in general choose among several possible successor states in $\delta(q, a)$.

Fairness assumptions are needed to give the learning algorithm sufficient opportunities to detect problems and try better policies (Francez, 1986). Intuitively, in both choice points, the choice should be independent of what the policy and working memory say about states other than the current state. This intuition is related to the Markov assumption, or *independence of path* assumption (Sutton and Barto, 1998). Below, we formalize this intuition as a fairness notion for the operational semantics of Section 3.2.2.

We say that a run $\mathcal{R}$ is *fair* if for each configuration c that occurs infinitely often at non-terminal positions, for each $(a, q') \in opt(c)$, the following transition occurs infinitely often:

$$c \xrightarrow{a,\, q'} apply(c, a, q').$$

We say that a task T is *learnable under fairness* if all fair runs of T converge.

Remark 3.7. There is always a fair run for any task, as follows. For each possible configuration c, we could conceptually order the set $opt(c)$. During a run, we could also keep track for each occurrence $i \in \mathbb{N}_0$ of a configuration c

[4] If a branching configuration c is terminal in a trial, c can not influence the action of its contained state because there is no subsequent transition anymore.

how many times we have already seen configuration c in the run, excluding the current occurrence; we denote this number as $count(i)$.

We begin the first trial with a random start configuration c_1, i.e., we choose a random start state and a random policy. We next choose the option $(a_1, q_2) \in opt(c_1)$ with the first ordinal in the now ordered set $opt(c_1)$. Now, for all the subsequent occurrences i of a configuration c in the run, we choose the option with ordinal $j = (count(i) \bmod |opt(c)|) + 1$ in the set $opt(c)$. So, if a configuration occurs infinitely often at non-terminal positions then we continually rotate through all its options. Naturally, trials end at the first occurrence of reward, and then we choose another start state; taking care to use all start states infinitely often. □

4 Results

The cycle-detection learning algorithm formalized in Section 3.2.2 continually marks the encountered states as visited. At the end of trials, i.e., after obtaining reward, each state is again marked as unvisited. If the algorithm encounters a state q that is already visited within the same trial, the algorithm proposes to generate a new action for q. Intuitively, if the same state q is encountered in the same trial, the agent might be running around in cycles and some new action should be tried for q to escape from the cycle. It is important to avoid cycles if we want to achieve an eventual upper bound on the length of a trial, i.e., an upper bound on the time it takes to reach reward from a given start state.

Repeatedly trying a new action for revisited states might eventually lead to reward, and thereby terminate the trial. In this learning process, the nondeterminism of the task can be both helpful and hindering: nondeterminism is helpful if transitions choose successor states that are closer to reward, but nondeterminism is hindering if transitions choose successor states that are further from reward or might lead to a cycle. Still, on some suitable tasks, like reducible tasks, the actions that are randomly tried upon revisits might eventually globally form a policy that will never get trapped in a cycle ever again (see Theorem 4.1 below).

The outline of this section is as follows. In Section 4.1, we present a sufficient condition for tasks to be learnable under fairness. In Section 4.2 we discuss how a simulator could detect that convergence has occurred in a fair run. In Section 4.3 we present necessary conditions for tasks to be learnable under fairness.

4.1 Sufficient Condition for Convergence

Intuitively, if a task is reducible then we might be able to obtain a policy that on each start state leads to reward without revisiting states in the same trial. As long as revisits occur, we keep searching for the acyclic flow of states that is implied by reducibility. We can imagine that states near the goal states, i.e., near immediate reward, tend to more quickly settle on an action that leads to reward. Subsequently, states that are farther removed from immediate reward can be reduced to states near goal states, and this growth process propagates through the entire state space. This intuition is confirmed by the following convergence result:

Theorem 4.1. All reducible tasks are learnable under fairness.

Proof. Let $T = (Q, Q_0, A, \mathit{rewards}, \delta)$ be a reducible task. Let $\mathcal{R}$ be a fair run on T. We show convergence of $\mathcal{R}$. In Part 1 of the proof, we show that all trials in $\mathcal{R}$ terminate (with reward). In Part 2, we show that eventually all states become stable in $\mathcal{R}$.

Part 1: Trials terminate. Let

$$L_1, L_2, \ldots$$

be the reducibility layers for T as defined in Section 3.1, where $L_1 = \mathit{goals}(T)$. Let $\mathcal{C}$ be a trial in $\mathcal{R}$. To show finiteness of $\mathcal{C}$, and thus termination of $\mathcal{C}$, we show by induction on $i = 1, 2, \ldots$ that the states in L_i occur only finitely many times in $\mathcal{C}$. Because T is reducible, there is an index $n \in \mathbb{N}_0$ for which $L_n = L_{n+1} = Q$, and therefore our inductive proof shows that every state only occurs a finite number of times in the trial $\mathcal{C}$; hence, $\mathcal{C}$ is finite.

Before we continue, we recall that a state q is marked as visited after its first occurrence in the trial; any occurrence of q after its first occurrence is therefore in a branching configuration.

Base case. Let $q \in L_1 = \mathit{goals}(T)$. Towards a contradiction, suppose q occurs infinitely often in trial $\mathcal{C}$, making $\mathcal{C}$ infinite. Because there are only a finite number of possible configurations, there is a configuration c containing q that occurs infinitely often in $\mathcal{C}$ at non-terminal positions (because the trial is now infinite). Configuration c is branching because it occurs more than once.[5] By definition of $q \in \mathit{goals}(T)$, there is an action $a \in A$ such that $(q, a) \in \mathit{rewards}$. Since always $\delta(q, a) \neq \emptyset$, we can choose some $q' \in \delta(q, a)$. We have $(a, q') \in \mathit{opt}(c)$ because c is branching. By fairness, the following transition must occur infinitely often in the trial:

$$c \xrightarrow{a,\, q'} \mathit{apply}(c, a, q').$$

But this transition is a reward transition, so the trial would have already ended at the first occurrence of this transition. Hence q can not occur infinitely many times; this is the desired contradiction.

Inductive step. Let $i \geq 2$, and let us assume that states in L_{i-1} occur only finitely many times in $\mathcal{C}$. Let $q \in L_i \setminus L_{i-1}$. By definition of L_i, there is some action $a \in A$ such that $\delta(q, a) \subseteq L_{i-1}$. Towards a contradiction, suppose q occurs infinitely often in $\mathcal{C}$, making $\mathcal{C}$ infinite. Like in the base case, there must be a branching configuration c containing q that occurs infinitely often in the trial (at non-terminal positions). Since always $\delta(q, a) \neq \emptyset$, we can choose some $q' \in \delta(q, a) \subseteq L_{i-1}$. We have $(a, q') \in \mathit{opt}(c)$ because c is branching. By fairness, the following transition must occur infinitely often in the trial:

$$c \xrightarrow{a,\, q'} \mathit{apply}(c, a, q').$$

But then q' would appear infinitely often in trial $\mathcal{C}$. This is the desired contradiction, because the induction hypothesis says that all states in L_{i-1} (including q') occur only finitely many times in $\mathcal{C}$.

[5] To see this, take for instance the second occurrence of c in the trial. That occurrence represents a revisit to q, so q is in the working memory set of c.

Part 2: Stability of states. We now show that all states eventually become stable in the fair $\mathcal{R}$. Let

$$L_1, L_2, \ldots$$

again be the reducibility layers for T as above, where $L_1 = \mathit{goals}(T)$. We show by induction on $i = 1, 2, \ldots$ that states in L_i become stable in $\mathcal{R}$. Since T is reducible, there is an index $n \in \mathbb{N}_0$ such that $L_n = L_{n+1} = Q$, so our inductive proof shows that all states eventually become stable.

Before we continue, we recall that Part 1 of the proof has shown that all trials are finite. So, whenever we say that a configuration occurs infinitely often in the run, this means that the configuration occurs in infinitely many trials. Similarly, if a transition occurs infinitely often in the run, this means that the transition occurs in infinitely many trials.

Base case. Let $q \in L_1 = \mathit{goals}(T)$. Towards a contradiction, suppose q would not become stable. This means that there are infinitely many non-terminal occurrences of branching configurations containing q.[6] Because there are only finitely many possible configurations, there must be a branching configuration c containing q that occurs infinitely often at non-terminal positions.

By definition of $q \in \mathit{goals}(T)$, there is an action $a \in A$ such that $(q, a) \in \mathit{rewards}$. Since always $\delta(q, a) \neq \emptyset$, we can choose some $q' \in \delta(q, a)$. We have $(a, q') \in \mathit{opt}(c)$ because c is branching. By fairness, the following transition t must occur infinitely often in the run:

$$c \xrightarrow{a,\, q'} \mathit{apply}(c, a, q').$$

Transition t is a reward transition because $(q, a) \in \mathit{rewards}$. Let j be the index of a trial containing transition t; this implies that t is the last transition of trial j. We now show that any non-terminal occurrences of q after trial j must be in a non-branching configuration. Hence, q becomes stable; this is the desired contradiction.

Consider the first trial index k after j in which q occurs again at a non-terminal position. Let configuration $c_1 = (q_1, \pi_1, w_1)$ with $q_1 = q$ be the first occurrence of q in trial k (which is at a non-terminal position). Note that $\pi_1(q) = a$ because *(i)* trial j ends with the assignment of action a to q (through transition t), and *(ii)* the trials between j and k could not have modified the action of q. Further, configuration c_1 is not branching because q is not yet flagged as visited at its first occurrence in trial k. This means that at any occurrence of c_1, trial k must select an option $(a, q'') \in \mathit{opt}(c_1)$, with action a and $q'' \in \delta(q, a)$, and perform the corresponding transition t':

$$c_1 \xrightarrow{a,\, q''} \mathit{apply}(c_1, a, q'').$$

Again, since $(q, a) \in \mathit{rewards}$, trial k ends directly after transition t'; no branching configuration containing q can occur in trial k at a non-terminal position.[7] This reasoning can now be repeated for all following trials to see that there are no more non-terminal occurrences of branching configurations containing q.

[6] For completeness, we recall that if q would occur only a finite number of times in the run then we can immediately see in the definition of stability that q becomes stable.

[7] Although it is possible that q is directly revisited from itself, it does not matter whether the terminal configuration of the trial is branching or not.

Inductive step. Let $i \geq 2$. We assume for each $q' \in L_{i-1}$ that q' eventually becomes stable. Now, let $q \in L_i \setminus L_{i-1}$. By definition of L_i, there is an action $a \in A$ such that $\delta(q, a) \subseteq L_{i-1}$. Towards a contradiction, suppose that q does not become stable. Our aim is to show that now also at least one $q' \in \delta(q, a)$ does not become stable, which would contradict the induction hypothesis.

Regarding terminology, we say that a chain is a (q, a)-*chain* if *(i)* the chain contains only non-reward transitions and *(ii)* the chain has the following desired form:

$$c_1 \xrightarrow{a_1,\, q_2} c_2 \xrightarrow{a_2,\, q_3} \dots \xrightarrow{a_{n-1},\, q_n} c_n,$$

denoting $c_j = (q_j, \pi_j, w_j)$ for each $j \in \{1, \dots, n\}$, where $q_1 = q_n = q$ and $a_1 = a$. Note that such a chain starts and ends with an occurrence of q, so q is revisited in the chain. Moreover, the first transition performs the action a from above. Next, we say that a trial is a (q, a)-*trial* if the trial contains a (q, a)-chain. In principle, each (q, a)-trial could embed a different (q, a)-chain.

To see that there are infinitely many occurrences of (q, a)-trials in $\mathcal{R}$, we distinguish between the following two cases.

- Suppose that in $\mathcal{R}$ there are infinitely many occurrences of trials that end with a policy π where $\pi(q) = a$, i.e., action a is assigned to q. Let j be the index of such a trial occurrence. Because by assumption q does not become stable, we can consider the first trial index k after j in which q occurs in a branching configuration at a non-terminal position. Note that trials between trial j and trial k do not modify the action of q. Now, the first occurrence of q in trial k is always non-branching, and thus we perform action a there. The subsequence in trial k starting at the first occurrence of q and ending at some branching configuration of q at a non-terminal position, is a (q, a)-chain: the chain starts and ends with q, its first transition performs action a, and it contains only non-reward transitions because it ends at a non-terminal position. Hence, trial k is a (q, a)-trial.

- Conversely, suppose that in $\mathcal{R}$ there are only finitely many occurrences of trials that end with a policy π where $\pi(q) = a$. Let $\mathcal{R}'$ be an (infinite) suffix of $\mathcal{R}$ in which no trial ends with action a assigned to q. Because by assumption q does not become stable, and because the number of possible configurations is finite, there is a branching configuration c containing q that occurs infinitely often at non-terminal positions in $\mathcal{R}'$. Choose some $q' \in \delta(q, a)$. We have $(a, q') \in \mathit{opt}(c)$ because c is branching. By fairness, the following transition t occurs infinitely often in $\mathcal{R}'$:

  $$c \xrightarrow{a,\, q'} \mathit{apply}(c, a, q').$$

  Let j be the index of a trial occurrence in $\mathcal{R}'$ that contains transition t; there are infinitely many such indexes because all trials in $\mathcal{R}$ are finite (see Part 1 of the proof). Since transition t attaches action a to q, we know by definition of $\mathcal{R}'$ that any occurrence of t in trial j is followed by at least one other transition from state q that attaches an action b to q with $b \neq a$; this implies that after each occurrence of transition t in trial j there is a branching configuration of q at a non-terminal position. In trial

j, a subsequence starting at any occurrence of t and ending with the first subsequent branching configuration of q at a non-terminal position, is a (q, a)-chain: the chain starts and ends with q, its first transition performs action a, and the chain contains only non-reward transitions because it ends at a non-terminal position. Hence, trial j is a (q, a)-trial.

We have seen above that there are infinitely many occurrences of (q, a)-trials in $\mathcal{R}$. Because there are only a finite number of possible configurations, there is a configuration c containing q that is used in infinitely many occurrences of (q, a)-trials as the last configuration of a (q, a)-chain. Note that c occurs infinitely often at non-terminal positions since (q, a)-chains contain no reward transitions.

Next, we can choose from some occurrence of a (q, a)-trial in the run some (q, a)-chain $\mathcal{C}$ where in particular the last configuration of $\mathcal{C}$ is the configuration c. Formally, we write $\mathcal{C}$ as

$$c_1 \xrightarrow{a_1,\, q_2} c_2 \xrightarrow{a_2,\, q_3} \dots \xrightarrow{a_{n-1},\, q_n} c_n,$$

where $c_n = c$, and denoting $c_j = (q_j, \pi_j, w_j)$ for each $j \in \{1, \dots, n\}$, where $q_1 = q_n = q$ and $a_1 = a$. We recall that all transitions of $\mathcal{C}$ are non-reward transitions. Note that $2 < n$: we have $q \neq q_2$ because $q \notin L_{i-1}$ and $q_2 \in \delta(q, a) \subseteq L_{i-1}$.

In chain $\mathcal{C}$, we have certainly marked state q as visited after its first occurrence, causing configuration c_n to be branching. This implies $(a, q_2) \in \mathit{opt}(c_n)$, where (a, q_2) is the same option as taken by the first transition of $\mathcal{C}$, since $a_1 = a$. Also, since $2 < n$, we have certainly marked state q_2 as visited after its first occurrence in $\mathcal{C}$; this implies $q_2 \in w_n$. Next, since the configuration $c_n = c$ occurs infinitely often at non-terminal positions (see above), the following transition also occurs infinitely often by fairness:

$$c_n \xrightarrow{a,\, q_2} c_{n+1},$$

where $c_{n+1} = (q_{n+1}, \pi_{n+1}, w_{n+1}) = \mathit{apply}(c_n, a, q_2)$. Because $q_{n+1} = q_2$ and $q_2 \in w_n \subseteq w_{n+1}$, configuration c_{n+1} is branching. Moreover, we know that $(q, a) \notin \mathit{rewards}$ since no transition of $\mathcal{C}$ is a reward transition, including the first transition. So, the branching configuration c_{n+1} occurs infinitely often at non-terminal positions. Hence, q_2 would not become stable. Yet, $q_2 \in \delta(q, a) \subseteq L_{i-1}$, and the induction hypothesis on L_{i-1} says that q_2 does become stable; this is the desired contradiction. □

Remark 4.2. By Theorem 4.1, the trials in a fair run on a reducible task eventually contain a number of non-terminal configurations that is at most the number of states; otherwise at least one state would never become stable.[8] So, we get a relatively good eventual upper bound on trial length. However, Theorem 4.1 provides no information on the waiting time before that upper bound will emerge, because that waiting time strongly depends on the choices made by the run regarding start states of trials, tried actions, and successor states (see also Section 6).

[8] If there would be infinitely many trials that contain more non-terminal configurations than states, then in infinitely many trials there is a revisit to a state (in a branching configuration) on a non-terminal position. Since there are only finitely many states, there would be at least one state q that in infinitely many trials occurs in a branching configuration on a non-terminal position; this state q does not become stable by definition.

Because we seek a policy that avoids revisits to states in the same trial, an important intuition implied by Theorem 4.1 is that for reducible tasks eventually the trials of a run follow paths without cycles through the state space. The followed paths are still influenced by nondeterminism, but they never contain a cycle. Also, a path followed in a trial is not necessarily the shortest possible path to reward, because the discovery of paths depends on experience, i.e., on the order in which actions were tried during the learning process. The experience dependence was experimentally observed, e.g. by Frémaux et al. (2013). □

Remark 4.3. The order in which states become stable in a fair run does not necessarily have to follow the order of the reducibility layers of Section 3.1. In general, it seems possible that some states that are farther removed from goal states could become stable faster than some states nearer to goal states; but, to become stable, the farther removed states probably should first have some stable strategy to the goal states.

To see that simulations do not exactly follow the inductive reasoning of the proof of Theorem 4.1, one could compare, in the later Section 5, the canonical policy implied by reducibility in Figure 5.2 with an actual final policy in Figure 5.4. □

The following example illustrates the necessity of the fairness assumption in Theorem 4.1. So, although the convergence result for reducible tasks appears natural, the example reveals that subtle notions, like the fairness assumption, should be taken into account to understand learning.

Example 4.4. Consider again the task T from Example 3.2, that is also visualized in Figure 3.1. In the following, for ease of notation, we will denote configurations as triples (x, y, Z), where x is the current state; y is the action assigned by the policy to the specific state 1, with action a assigned to all other states; and Z is the set of visited states as before.

Consider now the following trial $\mathcal{C}_a$ where the initial policy has assigned action a to all states, including the start state 1:

$$(1, a, \{\ \}) \xrightarrow{a,1} (1, a, \{1\}) \xrightarrow{b,2} (2, b, \{1\}) \xrightarrow{a,3} (3, b, \{1, 2\}) \xrightarrow{a,3} (3, b, \{1, 2, 3\}).$$

This is indeed a valid trial because the last transition is a reward transition. Note also that a revisit to state 1 occurs in the first transition. The configuration $(1, a, \{1\})$ is thus branching, which implies that the option $(b, 2)$ may be chosen there. At the end of trial $\mathcal{C}_a$, action b is assigned to state 1 and action a is assigned to the other states.

Consider also the following trial $\mathcal{C}_b$ where the initial policy has assigned action b to state 1 and a to all other states:

$$(1, b, \{\ \}) \xrightarrow{b,2} (2, b, \{1\}) \xrightarrow{a,1} (1, b, \{1, 2\}) \xrightarrow{a,3} (3, a, \{1, 2\}) \xrightarrow{a,3} (3, a, \{1, 2, 3\}).$$

The last transition is again a reward transition. Note that a revisit occurs to state 1 in the second transition. The configuration $(1, b, \{1, 2\})$ is therefore branching, which implies that the option $(a, 3)$ may be chosen there. At the end of trial $\mathcal{C}_b$, action a is assigned to all states, including state 1.

Now, let $\mathcal{R}$ be the run that alternates between trials $\mathcal{C}_a$ and $\mathcal{C}_b$ and that starts with trial $\mathcal{C}_a$. The state 1 never becomes stable in $\mathcal{R}$ because we assign action a and action b to state 1 in an alternating fashion. So, run $\mathcal{R}$ does not

converge because there are infinitely many non-terminal occurrences of branching configurations containing state 1.

Although run $\mathcal{R}$ satisfies all requirements of a valid run, $\mathcal{R}$ is not fair. For example, although the configuration $(1, b, \{1, 2\})$ occurs infinitely often (due to repeating trial $\mathcal{C}_b$), this configuration is never extended with the valid option $(b, 2)$ that could propagate revisits of state 1 to revisits of state 2 in the same trial; revisits to state 2 could force state 2 to use the other action b, which in turn could aid state 1 in becoming stable.

In conclusion, because task T is reducible and yet the valid (but unfair) run $\mathcal{R}$ does not converge, we see that Theorem 4.1 does not hold in absence of fairness. □

4.2 Detecting the Final Policy

We refer to the lastly formed policy of a run as the final policy. For an increased understanding of what convergence means, it appears interesting to say something about the form of the final policy. In particular, we would like to understand what kind of paths are generated by the final policy. As an additional benefit, recognizing the form of the final policy allows us to detect the convergence point in a simulation.[9]

We syntactically characterize the final policy in Theorem 4.5. In general, verifying the syntactical property of the final policy requires access to the entire set of task states. In this subsection, we do not require that tasks are reducible.

We first introduce the two key parts of the syntactical characterization, namely, the so-called *forward* and *backward* sets of states induced by a policy. As we will see below, the syntactical property says that the forward set should be contained in the backward set.

Forward and Backward Let $T = (Q, Q_0, A, \mathit{rewards}, \delta)$ be a task that is learnable under fairness. To make the notations below easier to read, we omit the symbol T from them. It will always be clear from the context which task is meant.

Let $\pi : Q \to A$ be a policy, i.e., each $q \in Q$ is assigned an action from A. First, we define

$$\mathit{ground}(\pi) = \{q \in \mathit{goals}(T) \mid (q, \pi(q)) \in \mathit{rewards}\};$$

this is the set of all goal states that are assigned a rewarding action by the policy. Next, we define two sets $\mathit{forward}(\pi) \subseteq Q$ and $\mathit{backward}(\pi) \subseteq Q$, as follows. For the set $\mathit{forward}(\pi)$, we consider the infinite sequence $F_1, F_2, \ldots$ of sets, where $F_1 = Q_0$ and for each $i \geq 2$,

$$F_i = F_{i-1} \cup \bigcup_{q \in F_{i-1} \setminus \mathit{ground}(\pi)} \delta(q, \pi(q)).$$

We define $\mathit{forward}(\pi) = \bigcup_{i \in \mathbb{N}_0} F_i$. Note that $\mathit{forward}(\pi) \subseteq Q$. Intuitively, the set $\mathit{forward}(\pi)$ contains all states that are reachable from the start states by

[9] Precise convergence detection is possible because our framework does not model reward numerically and thus there are no numerical instability issues near convergence. The convergence detection enables some of the simulation experiments in Section 5.

following the policy. In the definition of F_i with $i \geq 2$, we remove $\mathit{ground}(\pi)$ from the extending states because we only want to add states to $\mathit{forward}(\pi)$ that can occur at non-terminal positions of trials.[10]

For the set $\mathit{backward}(\pi)$, we consider the infinite sequence B_1, B_2, ... of sets, where $B_1 = \mathit{ground}(\pi)$ and for each $i \geq 2$,

$$B_i = B_{i-1} \cup \{q \in Q \mid \delta(q, \pi(q)) \subseteq B_{i-1}\}.$$

We define $\mathit{backward}(\pi) = \bigcup_{i \in \mathbb{N}_0} B_i$. Note that $\mathit{backward}(\pi) \subseteq Q$. Intuitively, $\mathit{backward}(\pi)$ is the set of all states that are reduced to the goal states in $\mathit{ground}(\pi)$ by the policy.

For completeness, we remark that the infinite sequences F_1, F_2, ..., and B_1, B_2, ..., each have a fixpoint because Q is finite.

Final Policy We formalize the final policy. Let T be a task that is learnable under fairness. Let $\mathcal{R}$ be a fair run on T, which implies that $\mathcal{R}$ converges. We define the *convergence-trial* of $\mathcal{R}$ as the smallest trial index i for which the following holds: trial i terminates and after trial i there are no more branching configurations at non-terminal positions.[11] This implies that after trial i the policy can not change anymore, because to change the action assigned to a state q, the state q would have to occur again in branching configuration at a non-terminal position. We define the *final policy* of $\mathcal{R}$ to be the policy at the end of the convergence-trial. In principle, different converging runs can have different final policies.

Now, we can recognize the final policy with the following property, that intuitively says that any states reachable by the policy are also safely reduced by the policy to reward:

Theorem 4.5. Let T be a task that is learnable under fairness.[12] Let $\mathcal{R}$ be a converging fair run of T. A policy π occurring in run $\mathcal{R}$ at the end of a trial is the final policy of $\mathcal{R}$ if and only if

$$\mathit{forward}(\pi) \subseteq \mathit{backward}(\pi).$$

Proof. We show in two separate parts that $\mathit{forward}(\pi) \subseteq \mathit{backward}(\pi)$ is *(i)* a sufficient and *(ii)* a necessary condition for π to be the final policy of run $\mathcal{R}$.

Part 1: Sufficient condition. Let π be a policy occurring in run $\mathcal{R}$ at the end of a trial. Assume that $\mathit{forward}(\pi) \subseteq \mathit{backward}(\pi)$. We show that π is the final policy of $\mathcal{R}$.

Concretely, we show that any trial starting with policy π will *(i)* use π in all its configurations, including the terminal configuration; and, *(ii)*, does not contain branching configurations at non-terminal positions. This implies that the first trial ending with π is the convergence-trial, so π is the final policy.

Let $\mathcal{C}$ be a trial in $\mathcal{R}$ that begins with policy π. We explicitly denote trial $\mathcal{C}$ as the following finite chain of transitions:

$$c_1 \xrightarrow{a_1, q_2} \ldots \xrightarrow{a_{n-1}, q_n} c_n.$$

[10] Possibly, some states directly reachable from $\mathit{ground}(\pi)$ are still in $\mathit{forward}(\pi)$ because those states are also reachable from states outside $\mathit{ground}(\pi)$.

[11] With this definition of convergence-trial, a run converges if and only if the run contains a convergence-trial.

[12] In contrast to Theorem 4.1, we do not require that T is reducible.

For each $i \in \{1, \ldots, n\}$, we denote $c_i = (q_i, \pi_i, w_i)$. Let $F_1, F_2, \ldots$ be the infinite sequence of sets previously defined for $\mathit{forward}(\pi)$. We show by induction on $i = 1, \ldots, n-1$ that

(a) $\pi_i = \pi$;

(b) $q_i \in F_i$;

(c) c_i is non-branching.

At the end of the induction, we can also see that $\pi_n = \pi$: first, we have $\pi_n = \pi_{n-1}$ because configuration c_{n-1} is non-branching by property (c);[13] second, $\pi_{n-1} = \pi$ by property (a).

Base case. Let $i = 1$. For property (a), we have $\pi_1 = \pi$ because the trial starts with policy π. For property (b), we see that $q_1 \in Q_0 = F_1$. For property (c), we know that c_1 is non-branching because the first configuration in a trial still has an empty working memory of visited states.

Inductive step. Let $i \geq 2$, with $i \leq n-1$. Assume that the induction properties are satisfied for the configurations $c_1, \ldots, c_{i-1}$. We now show that the properties are also satisfied for c_i.

Property (a) By applying the induction hypothesis for property (c) to c_{i-1}, namely that c_{i-1} is non-branching, we know $\pi_i = \pi_{i-1}$. By subsequently applying the induction hypothesis for property (a) to c_{i-1}, namely $\pi_{i-1} = \pi$, we know $\pi_i = \pi$.

Property (b) To start, we note that $q_i \in \delta(q_{i-1}, \pi_{i-1}(q_{i-1}))$ because c_{i-1} is non-branching by the induction hypothesis for property (c). By subsequently applying the induction hypothesis for property (a) to c_{i-1}, namely $\pi_{i-1} = \pi$, we know $q_i \in \delta(q_{i-1}, \pi(q_{i-1}))$. Moreover, since $i-1 < i \leq n-1$, the transition $c_{i-1} \xrightarrow{a_{i-1},\, q_i} c_i$, where $a_{i-1} = \pi(q_{i-1})$, is a non-reward transition. Hence, $(q_{i-1}, \pi(q_{i-1})) \notin \mathit{rewards}$ and thus $q_{i-1} \notin \mathit{ground}(\pi)$.

Lastly, by applying the induction hypothesis for property (b) to c_{i-1}, we overall obtain that $q_{i-1} \in F_{i-1} \setminus \mathit{ground}(\pi)$. Combined with $q_i \in \delta(q_{i-1}, \pi(q_{i-1}))$, we see that $q_i \in F_i$.

Property (c) Towards a contradiction, suppose that configuration c_i is branching. This means that state q_i is revisited in c_i.[14] Let $V = \{q_1, \ldots, q_{i-1}\}$. Note that $q_i \in V$, which implies $V \neq \emptyset$. By applying the induction hypothesis for property (b) to configurations $c_1, \ldots, c_{i-1}$, we know that $V \subseteq \mathit{forward}(\pi)$. We now show that $V \cap \mathit{backward}(\pi) = \emptyset$, which would imply $\mathit{forward}(\pi) \not\subseteq \mathit{backward}(\pi)$; this is the desired contradiction.

Let $B_1, B_2, \ldots$ be the infinite sequence of sets defined for $\mathit{backward}(\pi)$ above. We show by induction on $j = 1, 2, \ldots$ that $V \cap B_j = \emptyset$, which then overall implies $V \cap \mathit{backward}(\pi) = \emptyset$.

[13] Because configuration c_{n-1} is non-branching, we have $\pi_n(q_{n-1}) = a_{n-1} = \pi_{n-1}(q_{n-1})$, which, combined with $\pi_n(r) = \pi_{n-1}(r)$ for each $r \in Q \setminus \{q_{n-1}\}$, gives $\pi_n = \pi_{n-1}$.

[14] Recall that, by definition of branching configuration, we have $q_i \in w_i$.

- Base case: $j = 1$. By definition, $B_1 = \mathit{ground}(\pi)$. Let $q \in V$. Let $k \in \{1, \ldots, i-1\}$ be the smallest index for which $q_k = q$, i.e., configuration c_k represents the first occurrence of q in the trial. By applying the outer induction hypothesis for properties (a) and (c) to c_k, we know that $a_k = \pi_k(q_k) = \pi(q)$. But since $k \leq i-1 < i \leq n-1$, we know that transition $c_k \xrightarrow{a_k,\, q_{k+1}} c_{k+1}$ is not a reward transition, implying $(q, \pi(q)) \notin \mathit{rewards}$. Hence, $q \notin \mathit{ground}(\pi)$, and overall $V \cap \mathit{ground}(\pi) = \emptyset$.
- Inductive step. Let $j \geq 2$. Assume $V \cap B_{j-1} = \emptyset$. Towards a contradiction, suppose $V \cap B_j \neq \emptyset$. Take some $q \in V \cap B_j$. If $q \in B_{j-1}$ then we would immediately have a contradiction with the induction hypothesis. Henceforth, suppose $q \in B_j \setminus B_{j-1}$, which, by definition of B_j, means $\delta(q, \pi(q)) \subseteq B_{j-1}$. We will now show that $V \cap \delta(q, \pi(q)) \neq \emptyset$, which would give $V \cap B_{j-1} \neq \emptyset$; this is the desired contradiction.

  Since $q \in V$, there is some smallest $k \in \{1, \ldots, i-1\}$ such that $q_k = q$. Using a similar reasoning as in the base case ($j = 1$), by applying the outer induction hypothesis for properties (a) and (c) to configuration c_k, we can see that $a_k = \pi(q_k)$. This implies $q_{k+1} \in \delta(q_k, \pi(q_k))$. As a last step, we show that $q_{k+1} \in V$, which gives $V \cap \delta(q_k, \pi(q_k)) \neq \emptyset$. We distinguish between the following cases:
  - If $k \leq i-2$ then $k+1 \leq i-1$, and surely $q_{k+1} \in V$ by definition of V.
  - If $k = i-1$ then we know $q_{k+1} = q_i \in V$ because configuration c_i revisits state q_i (see above).

Part 2: Necessary condition. Let π be the final policy of $\mathcal{R}$. We show that $\mathit{forward}(\pi) \subseteq \mathit{backward}(\pi)$. By definition of final policy, π is the policy at the end of the convergence-trial, whose trial index we denote as i. By definition of convergence-trial, after trial i there are no more branching configurations at non-terminal positions. Note in particular that the policy no longer changes after trial i.

Towards a contradiction, suppose that $\mathit{forward}(\pi) \not\subseteq \mathit{backward}(\pi)$. Let $V = \mathit{forward}(\pi) \setminus \mathit{backward}(\pi)$. Note that $V \neq \emptyset$. We show that there is a state $q \in V$ that occurs at least once in a branching configuration at a non-terminal position after the convergence-trial i; this would be the desired contradiction.

We provide an outline of the rest of the proof. The reasoning proceeds in two steps. First, we show for each $q \in V$ that $\delta(q, \pi(q)) \cap V \neq \emptyset$. This means that if we are inside set V, we have the option to stay longer inside V if we follow the policy π. Now, the second step of the reasoning is to show that we can stay arbitrarily long inside V even after the convergence-trial i, causing at least one state of V to occur in a branching configuration at a non-terminal position after trial i.

Step 1. Let $q \in V$. We show $\delta(q, \pi(q)) \cap V \neq \emptyset$. Towards a contradiction, suppose that $\delta(q, \pi(q)) \cap V = \emptyset$. Our strategy is to show that $\delta(q, \pi(q)) \subseteq \mathit{backward}(\pi)$, which, by definition of $\mathit{backward}(\pi)$, implies that there is some index $j \in \mathbb{N}_0$ such that $\delta(q, \pi(q)) \subseteq B_j$. Therefore $q \in B_{j+1} \subseteq \mathit{backward}(\pi)$.

But that is false because $q \in V = \mathit{forward}(\pi) \setminus \mathit{backward}(\pi)$; this is the desired contradiction.

We are left to show that $\delta(q, \pi(q)) \subseteq \mathit{backward}(\pi)$. First, we show $\delta(q, \pi(q)) \subseteq \mathit{forward}(\pi)$. By definition, $\mathit{forward}(\pi) = \bigcup_{j \in \mathbb{N}_0} F_j$. Since $q \in V \subseteq \mathit{forward}(\pi)$, there is some index $j \in \mathbb{N}_0$ such that $q \in F_j$. Moreover, since $q \notin \mathit{backward}(\pi)$ and $\mathit{ground}(\pi) \subseteq \mathit{backward}(\pi)$, we have $q \notin \mathit{ground}(\pi)$. Overall, $q \in F_j \setminus \mathit{ground}(\pi)$, which implies that $\delta(q, \pi(q)) \subseteq F_{j+1} \subseteq \mathit{forward}(\pi)$.

Now, we can complete the reasoning by combining $\delta(q, \pi(q)) \subseteq \mathit{forward}(\pi)$ with our assumption $\delta(q, \pi(q)) \cap V = \emptyset$, to see the following:

$$\begin{aligned}\delta(q, \pi(q)) &\subseteq \mathit{forward}(\pi) \setminus V \\ &= \mathit{forward}(\pi) \setminus (\mathit{forward}(\pi) \setminus \mathit{backward}(\pi)) \\ &= \mathit{forward}(\pi) \cap \mathit{backward}(\pi) \\ &\subseteq \mathit{backward}(\pi).\end{aligned}$$

Step 2. We now show that after convergence-trial i there is at least one occurrence of a branching configuration at a non-terminal position.

We first show that each $q \in \mathit{forward}(\pi)$ occurs infinitely often at non-terminal positions after trial i. Recall that $\mathit{forward}(\pi) = \bigcup_{j \in \mathbb{N}_0} F_j$. We show by induction on $j = 1, 2, \ldots$ that states in F_j occur infinitely often at non-terminal positions after trial i.

- Base case: $j = 1$. By definition, $F_1 = Q_0$. Because $\mathcal{R}$ is a valid run, each state of Q_0 is used in infinitely many trials as the start state, also after trial i (see Section 3.2.2).[15] Moreover, the first configuration of a trial is always at a non-terminal position because each trial contains at least one transition.

- Inductive step. Let $j \geq 2$. Assume that each state in F_{j-1} occurs infinitely often at non-terminal positions after trial i. Let $q \in F_j \setminus F_{j-1}$. This implies that there is some $q_{j-1} \in F_{j-1} \setminus \mathit{ground}(\pi)$ for which $q \in \delta(q_{j-1}, \pi(q_{j-1}))$.

  By applying the induction hypothesis, we know that q_{j-1} occurs infinitely often at non-terminal positions after trial i. Because there are only a finite number of possible configurations, there is a configuration c containing q_{j-1} that occurs infinitely often at non-terminal positions after trial i. Because trial i is the convergence-trial, we can make two observations about configuration c: first, c contains the final policy π because the policy no longer changes after trial i; and, second, c is non-branching because no branching configurations occur at non-terminal positions after trial i. These observations imply $(\pi(q_{j-1}), q) \in \mathit{opt}(c)$.

  Now, since c occurs infinitely often at non-terminal positions after trial i, the following transition t occurs infinitely often after trial i by fairness:

  $$c \xrightarrow{\pi(q_{j-1}),\, q} \mathit{apply}(c, \pi(q_{j-1}), q).$$

  Lastly, we know that $(q_{j-1}, \pi(q_{j-1})) \notin \mathit{rewards}$ because $q_{j-1} \notin \mathit{ground}(\pi)$, so transition t is a non-reward transition. Therefore state q occurs infinitely often at non-terminal positions after trial i.

[15] We also recall here that $\mathcal{R}$ is assumed to converge, which, by definition of convergence, implies that $\mathcal{R}$ is an infinite sequence of finite trials.

Now take some $q_1 \in V$. Since $V \subseteq \mathit{forward}(\pi)$, we know from above that q_1 occurs infinitely often at non-terminal positions after trial i. Because there are only finitely many possible configurations, there is a configuration c_1 containing q_1 that occurs infinitely often at non-terminal positions after trial i. After trial i, the policy no longer changes and only non-branching configurations may occur at non-terminal positions. So c_1 contains the final policy π and is non-branching. Moreover, since $q_1 \in V$, we know from Step 1 above that there is some $q_2 \in \delta(q_1, \pi(q_1)) \cap V$. Overall, we can see that $(\pi(q_1), q_2) \in \mathit{opt}(c_1)$. By fairness, the following transition t_1 occurs infinitely often after trial i:

$$c_1 \xrightarrow{\pi(q_1),\, q_2} c_2,$$

where $c_2 = \mathit{apply}(c_1, \pi(q_1), q_2)$. Because $q_1 \in V$ we have $q_1 \notin \mathit{ground}(\pi)$, so transition t_1 is a non-reward transition.[16] Therefore configuration c_2 occurs infinitely often at non-terminal positions after trial i. Denoting $c_2 = (q_2, \pi_2, w_2)$, note that $q_1 \in w_2$.

We now make a similar reasoning for c_2 as we did for c_1. Since $q_2 \in V$, we know again from Step 1 that there is some $q_3 \in \delta(q_2, \pi(q_2)) \cap V$. Configuration c_2 contains the final policy π because c_2 occurs after trial i, and c_2 is also non-branching because it occurs after trial i at a non-terminal position. Therefore $(\pi(q_2), q_3) \in \mathit{opt}(c_2)$. Now, since c_2 occurs infinitely often at non-terminal positions after trial i (see above), the following transition occurs infinitely often after trial i by fairness:

$$c_2 \xrightarrow{\pi(q_2),\, q_3} c_3,$$

where $c_3 = \mathit{apply}(c_2, \pi(q_2), q_3)$. This transition is also non-rewarding because $q_2 \in V$ implies $q_2 \notin \mathit{ground}(\pi)$. Denoting $c_3 = (q_3, \pi_3, w_3)$, note that $\{q_1, q_2\} \subseteq w_3$. We emphasize that more states of V are now marked as visited.

We can now complete the reasoning. By following the final policy π from a state in V, we can always stay inside V without reaching reward. So, the above procedure can be repeated $|V|$ times in total, to show the existence of a configuration $c = (q, \pi, w)$ with $q \in V$ that occurs infinitely often at non-terminal positions after trial i, and where $q \in w$. Configuration c is therefore branching, and thus the existence of c gives the desired contradiction, as explained at the beginning of Part 2 of the proof.[17] □

4.3 Necessary Conditions for Convergence

In Section 4.1, we have seen that reducibility is a sufficient property for tasks to be learnable under fairness (Theorem 4.1). In this subsection, we also show necessary properties for tasks to be learnable under fairness. This provides a first step towards characterizing the tasks that are learnable under fairness. Thinking about such a characterization is useful because it allows us to better understand the tasks that are not learnable under fairness.

We first introduce some auxiliary concepts. Let $T = (Q, Q_0, A, \mathit{rewards}, \delta)$ be a task. We say that there is a *(simple) path* from a state q to a state q' if

[16] If $q_1 \in \mathit{ground}(\pi)$ then $q_1 \in \mathit{backward}(\pi)$, which is false since $q_1 \in V$.

[17] We note the following for completeness. By repeatedly trying to stay inside V, we are not guaranteed to see all of V, but we still know that after at most $|V|$ transitions we have to revisit a state of V at a non-terminal position (after trial i); that first revisit forms the desired contradiction.

there is a sequence of actions $a_1, \ldots, a_n$, with possibly $n = 0$, and a sequence of states $q_1, \ldots, q_{n+1}$ such that

- $q_1 = q$;
- $q_{n+1} = q'$;
- $q_{i+1} \in \delta(q_i, a_i)$ for each $i \in \{1, \ldots, n\}$;
- $(q_i, a_i) \notin \mathit{rewards}$ for each $i \in \{1, \ldots, n\}$; and,
- $q_i \neq q_j$ for each $i, j \in \{1, \ldots, n+1\}$ with $i \neq j$.

We emphasize that the path does not contain reward transitions and does not repeat states. We also denote a path as

$$q_1 \xrightarrow{a_1} \ldots \xrightarrow{a_n} q_{n+1}.$$

Note that there is always a path from a state to itself, namely, the empty path where $n = 0$.

We say that a state q is *reachable* if there is a path from a start state $q_0 \in Q_0$ to q. Next, letting $V \subseteq Q$, we say that a state q *has a path to* V if there is a path from q to a state $q' \in V$. Note that if q has a path to V, it is not guaranteed that the action sequence of that path always ends in V because the transition function δ is nondeterministic.

We can now note the following necessary properties for tasks to be learnable under fairness:

Proposition 4.6. Tasks T that are learnable under fairness satisfy the following properties:

(a) The reachable states have a path to $\mathit{goals}(T)$.

(b) The start states are reducible to $\mathit{goals}(T)$.

The property (a) is related to the assumption that reward can be reached from every state, see e.g. (Bonet and Geffner, 2006).

Proof. We show the two properties separately. Denote $T = (Q, Q_0, A, \mathit{rewards}, \delta)$.

Property (a). Towards a contradiction, suppose that there is a reachable state q that has no path to $\mathit{goals}(T)$. Note that $q \notin \mathit{goals}(T)$ because otherwise the empty path from q to itself would be a path to $\mathit{goals}(T)$. Because q is reachable, there is a path

$$q_1 \xrightarrow{a_1} \ldots \xrightarrow{a_n} q_{n+1},$$

where $q_1 \in Q_0$ and $q_{n+1} = q$. We can consider a policy π where, for each $i \in \{1, \ldots, n\}$, we set $\pi(q_i) = a_i$; the other state-action mappings may be arbitrary.[18] Below we will consider a fair run $\mathcal{R}$ whose first trial is given q_1 as start state and π as the initial policy. First, we consider the following chain $\mathcal{C}$:

$$c_1 \xrightarrow{a_1,\, q_2} \ldots \xrightarrow{a_n,\, q_{n+1}} c_{n+1},$$

[18] Note that there are no conflicting assignments of actions to states because, by definition of path, each state q_i with $i \in \{1, \ldots, n+1\}$ occurs only once on the path.

where for each $i \in \{1, \ldots, n{+}1\}$ we define $c_i = (q_i, \pi, w_i)$ where $w_i = \{q_1, \ldots, q_{i-1}\}$.[19] Note that this chain is indeed valid: for each $i \in \{1, \ldots, n\}$, configuration c_i is (constructed to be) non-branching and therefore $(\pi(q_i), q_{i+1}) = (a_i, q_{i+1}) \in \mathit{opt}(c_i)$; this means that we do not modify the policy during the transition $c_i \xrightarrow{a_i,\, q_{i+1}} \mathit{apply}(c_i, a_i, q_{i+1})$, but we only mark q_i as visited, which gives $c_{i+1} = \mathit{apply}(c_i, a_i, q_{i+1})$. Note that configuration c_{n+1} contains the state q.

Consider a fair run $\mathcal{R}$ whose first trial starts with the chain $\mathcal{C}$. We show that the first trial never terminates; this is the desired contradiction because we had assumed that task T is learnable under fairness. For the first trial to terminate, the trial must extend chain $\mathcal{C}$ to a chain $\mathcal{C}'$ that terminates with a reward transition. But then the existence of $\mathcal{C}'$ would imply that there is a path from q to a state $q' \in \mathit{goals}(T)$, which is false.

Property (b). We show that the start states are reducible to $\mathit{goals}(T)$. Let $\mathcal{R}$ be a fair run on T, which implies that $\mathcal{R}$ converges. Because $\mathcal{R}$ converges, there is a final policy π, as defined in Section 4.2. Then by Theorem 4.5, we know that $\mathit{forward}(\pi) \subseteq \mathit{backward}(\pi)$. We have $Q_0 \subseteq \mathit{backward}(\pi)$ because $Q_0 \subseteq \mathit{forward}(\pi)$ by definition of $\mathit{forward}(\pi)$. Letting $\mathit{reduce}(T)$ be the set of states that are reducible to $\mathit{goals}(T)$, we show below that $\mathit{backward}(\pi) \subseteq \mathit{reduce}(T)$; this implies $Q_0 \subseteq \mathit{reduce}(T)$, as desired.

In Section 3.1, we have defined $\mathit{reduce}(T) = \bigcup_{i \in \mathbb{N}_0} L_i$ where $L_1 = \mathit{goals}(T)$ and for each $i \geq 2$,

$$L_i = L_{i-1} \cup \{q \in Q \mid \exists a \in A \text{ with } \delta(q, a) \subseteq L_{i-1}\}.$$

Also recall the definition $\mathit{backward}(\pi) = \bigcup_{i \in \mathbb{N}_0} B_i$ from Section 4.2. We show by induction on $i = 1, 2, \ldots$ that $B_i \subseteq L_i$.

- Base case: $i = 1$. By definition, $B_1 = \mathit{ground}(\pi) = \{q \in \mathit{goals}(T) \mid (q, \pi(q)) \in \mathit{rewards}\}$. Hence, $B_1 \subseteq \mathit{goals}(T) = L_1$.
- Inductive step. Let $i \geq 2$, and let us assume that $B_{i-1} \subseteq L_{i-1}$. Let $q \in B_i \setminus B_{i-1}$, which, by definition of B_i, implies $\delta(q, \pi(q)) \subseteq B_{i-1}$. By applying the induction hypothesis, we see $\delta(q, \pi(q)) \subseteq L_{i-1}$. This implies $q \in L_i$, and, overall, $B_i \subseteq L_i$.

This completes the proof. □

We recall from Theorem 4.1 that reducibility is a sufficient property for tasks to be learnable under fairness. Note that reducibility implies the necessary properties in Proposition 4.6 for tasks to be learnable under fairness. We now discuss the gap between these sufficient and necessary properties. Let $T = (Q, Q_0, A, \mathit{rewards}, \delta)$ be a task. Letting $\mathit{reduce}(T)$ be the set of states that can be reduced to $\mathit{goals}(T)$ (as in Section 3.1), we define

$$\mathit{unstable}(T) = Q \setminus \mathit{reduce}(T).$$

We call $\mathit{unstable}(T)$ the *unstable set* of T. Note that for each state $q \in \mathit{unstable}(T)$, for each action $a \in A$, we always have $\delta(q, a) \cap \mathit{unstable}(T) \neq \emptyset$.[20]

[19] In this notation, we interpret $\{q_1, \ldots, q_0\}$ as $\emptyset$.

[20] Indeed, if $\delta(q, a) \cap \mathit{unstable}(T) = \emptyset$ then $\delta(q, a) \subseteq \mathit{reduce}(T)$ and subsequently $q \in \mathit{reduce}(T)$, which is false.

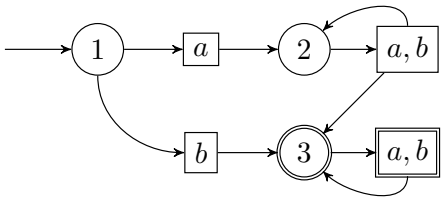

Figure 4.1: The task from Example 4.7, with unstable set $\{2\}$. The graphical notation is explained in Figure 3.1.

So, once we are inside $\mathit{unstable}(T)$, we can never reliably escape $\mathit{unstable}(T)$: escaping $\mathit{unstable}(T)$ depends on the nondeterministic choices regarding successor states. This intuition has also appeared in Part 2 of the proof of Theorem 4.5, but in that proof we were focusing on just a single fixed action for each state, as assigned by the final policy at hand.

The following example illustrates how a nonempty unstable set could prevent convergence. In particular, the example illustrates that the necessary properties of Proposition 4.6 are not sufficient for a task to be learnable under fairness.

Example 4.7. Consider the task $T = (Q, Q_0, A, \mathit{rewards}, \delta)$ defined as follows: $Q = \{1, 2, 3\}$; $Q_0 = \{1\}$; $A = \{a, b\}$; $\mathit{rewards} = \{(3, a), (3, b)\}$; and, regarding δ, we define

$$\begin{aligned} \delta(1, a) &= \{2\}, \\ \delta(1, b) &= \{3\}, \\ \delta(2, a) = \delta(2, b) &= \{2, 3\}, \\ \delta(3, a) = \delta(3, b) &= \{3\}. \end{aligned}$$

The task T is visualized in Figure 4.1. Note that $\mathit{reduce}(T) = \{1, 3\}$, giving

$$\mathit{unstable}(T) = \{2\}.$$

Note in particular that $\delta(2, a) \cap \mathit{unstable}(T) \neq \emptyset$ and $\delta(2, b) \cap \mathit{unstable}(T) \neq \emptyset$.

Task T satisfies the necessary properties of Proposition 4.6: *(i)* the reachable states, which are all states in this case, have a path to $\mathit{goals}(T)$, and *(ii)* the start state 1 is reducible to $\mathit{goals}(T)$. However, task T is not learnable under fairness, as we now illustrate. Consider a trial $\mathcal{C}$ of the following form: starting with an initial policy that assigns action a to state 1, we first go from state 1 to state 2; next, we stay at least two consecutive times in state 2; and, lastly, we proceed to state 3 and obtain reward there. Because there are no revisits to state 1 in trial $\mathcal{C}$, state 1 remains connected to action a. We now see that we can make a fair run $\mathcal{R}$ by repeating trials of the form of $\mathcal{C}$: state 1 is never revisited and stays connected to action a, and we keep revisiting state 2. This way, there are infinitely many branching configurations containing state 2 at non-terminal positions. So, run $\mathcal{R}$ does not converge. □

By looking at Example 4.7, it seems that runs would somehow have to learn to avoid entering the set $\mathit{unstable}(T)$. First, we define

$$\mathit{border}(T) = \{q \in \mathit{reduce}(T) \mid \exists a \in A \text{ with } \delta(q, a) \cap \mathit{unstable}(T) \neq \emptyset\}.$$

The set $\mathit{border}(T)$ contains those states that could enter $\mathit{unstable}(T)$; we call such states *border states*. Let $q \in \mathit{border}(T)$, and let V be the subset of $\mathit{unstable}(T)$ that is reachable from q. Now, one idea could be to demand for each $q_1 \in V$ that there is some $(q_2, a) \in V \times A$ such that $q_2 \in V$ is reachable from q_1 and $q \in \delta(q_2, a)$, i.e., there is some escape option to return from V back to border state q. This way, we can revisit that precise border state q in the same trial, so that under fairness we can choose a new action for q to avoid a future entrance into $\mathit{unstable}(T)$. The possibility to revisit border states in the same trial is exactly what is missing from Example 4.7. The characterization of tasks that are learnable under fairness might be a class of tasks that satisfy the necessary properties of Proposition 4.6 and that additionally specify assumptions on the transition function to ensure the possibility of revisits to border states. Illuminating the role of unstable sets in learning is an interesting avenue for further work (see Section 6).

5 Examples

Theorem 4.1 tells us that all reducible tasks are learnable under fairness, and Theorem 4.5 allows us to detect when the final policy has been formed. To illustrate these theorems, we now consider two examples of tasks that are reducible, in Section 5.1 and Section 5.2, respectively. Our aim is not to show practical efficiency of the cycle-detection learning algorithm, but rather to illustrate the theoretical insights. Indeed, because the considered examples are reducible, they are learnable under fairness by Theorem 4.1. Next, aided by Theorem 4.5, we can experimentally measure how long it takes for the learning process to convergence. In Section 6, for further work, we identify aspects where the learning algorithm could be improved to become more suitable for practice.

5.1 Grid Navigation Tasks

Our first example is a navigation task in a grid world (Sutton and Barto, 1998; Potjans et al., 2011). In such a task, it is intuitive to imagine how paths are formed and what they mean. Below we formalize a possible version of such a grid task.

The grid is represented along two axes, the X- and Y-axis. At any time, the agent is inside only one grid cell (x, y). We let the state set Q be a subset of $\mathbb{N} \times \mathbb{N}$. Let $\mathit{Goals} \subseteq Q$ be a subset of cells, called *goal cells*. The agent could apply the following actions to each grid cell: finish, left, right, up, down, left-up, left-down, right-up, and right-down. Let A_{grid} denote the set containing these actions. The finish action is a non-movement action that gives immediate reward when applied to a goal cell. The finish action intuitively says that the agent believes it has reached a goal cell and claims to be finished. Activating the finish action in any non-goal cell will just leave the agent in that cell without reward. The actions other than finish will be referred to as movement actions.

For every movement action a, there is noise from the environment. We formalize this with a noise offset function that maps each movement action to the possible relative movements that it can cause. For example, $\mathit{offset}(\text{left}) = \{(-1, 0), (-1, -1), (-1, 1)\}$, and $\mathit{offset}(\text{left-up}) = \{(-1, 1), (-1, 0), (0, 1)\}$. Intuitively, noise adds one left-rotating option and one right-rotating option to

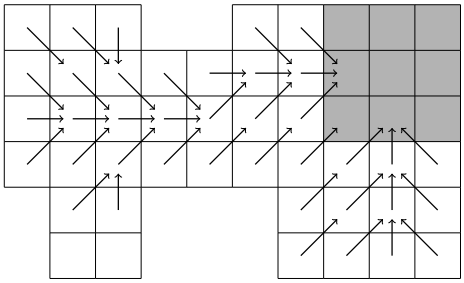

Figure 5.1: A grid navigation task. The shaded cells are goal cells, where the finish action should be executed to obtain reward. Reducibility of cells is illustrated by the arrows: an arrow leaving a cell represents an action that the cell could choose to come closer to a goal cell. We note that each arrow displays only the main intended direction of the action, and that in principle the noise offsets could occur during execution. In this example, not all cells are reducible. For the non-reducible cells, for each applied movement action a, either *(i)* function δ defines a to stay stationary for the cell, or *(ii)* a could lead to another non-reducible cell. Naturally, if all non-reducible cells are removed then the resulting task is reducible.

the main intended direction. The offsets of the other movement actions can be similarly defined (see Appendix A). For uniformity we define $\mathit{offset}(\text{finish}) = \{(0,0)\}$.

For a cell $(x, y) \in \mathbb{N}\times\mathbb{N}$ and an action $a \in A_{\text{grid}}$ we define the set $\mathit{move}((x, y), a)$ of child-cells that result from the application of the offsets of a to (x, y). Formally, we define

$$\mathit{move}((x, y), a) = \{(x + u,\, y + v) \mid (u, v) \in \mathit{offset}(a)\}.$$

Now, we say that a task $T = (Q, Q_0, A, \mathit{rewards}, \delta)$ is a *grid navigation task* if $Q \subseteq \mathbb{N}\times\mathbb{N}$, $Q_0 \subseteq Q$, $A = A_{\text{grid}}$, there exists some nonempty subset $\mathit{Goals} \subseteq Q$ such that

$$\mathit{rewards} = \{(q, \text{finish}) \mid q \in \mathit{Goals}\},$$

and for each $(q, a) \in Q \times A$ we have

$$\delta(q, a) = \begin{cases} \mathit{move}(q, a) & \text{if } \mathit{move}(q, a) \subseteq Q \\ \{q\} & \text{otherwise.} \end{cases}$$

Note that we only perform a movement action if the set of child-cells is fully contained in the set of possible grid cells; otherwise the agent remains stationary.[21]

The assumption of reducibility can additionally be imposed on grid navigation tasks. In Figure 5.1, we visualize a grid navigation task that, for illustrative purposes, is only partially reducible.

We now discuss two simulation experiments that we have performed on such grid navigation tasks. Some details of the experiments can be found in Appendix A. Let $\mathbb{N}_0$ denote the set of natural numbers without zero.

[21] This restriction results in policies that have sufficiently intuitive visualizations; see Figures 5.1, 5.2, and 5.4.

Convergence-Trial Index Our first experiment measures the convergence-trial index. First we discuss the general setup of the experiment. We recall from Section 4.2 that the convergence-trial index of a fair run is the first trial index where the final policy occurs at the end. For a given task T, we can simulate fair runs, and we stop each run when the final policy is detected through the characterization of Theorem 4.5; we remember the convergence-trial index. Each run is started with a random policy, where each state is assigned a random action by uniformly sampling the available actions. Fairness depends on the mechanism for choosing among successor states, which is also based on a uniform distribution (see Appendix A). Interestingly, we do not have to simulate infinite runs because we always eventually detect the final policy; when we stop the simulation after finite time, the finite run may be thought of as a prefix of an infinite fair run.[22]

In principle, there is no upper bound on the convergence-trial index in the simulation. Fortunately, our experiments demonstrate that for the studied reducible tasks, there appears to be a number $i \in \mathbb{N}_0$ such that the large majority of simulated runs has a convergence-trial index below i. Possibly, there are outliers with a very large convergence-trial index, although such outliers are relatively few in number. We can exclude outliers from a list of numbers by considering a p-quantile with $0 < p < 1$.[23]

We wanted to see if the convergence-trial index depends on the distance between the start cells and the goal cells. For this purpose, we have considered grid tasks with the form of a corridor, as shown in Figure 5.2. There is only one start cell. The parameter that we can change is the distance from the start cell to the patch of goal cells. All other aspects remain fixed, including the width of the corridor and the number and the location of the goal cells. The arrows in Figure 5.2 show reducibility of this kind of task. For a number $l \in \mathbb{N}_0$, we define the l-corridor as the corridor navigation task where the distance between the start cell and the goal cells is equal to l.

Now, for some lengths l, we have simulated runs on the l-corridor. For each l-corridor separately, we have simulated 400 runs and we have computed the 0.9-quantile for the measured convergence-trial indexes; this gives an empirical upper bound on the convergence-trial index for the l-corridor. Figure 5.3a shows the quantiles plotted against the corridor lengths. We observe that longer corridors require more time to learn. This is probably because at each application of a movement action to a cell, there could be multiple successor cells due to nondeterminism. Intuitively, the nondeterminism causes a drift away from any straight path to the goal cells. The policy has to learn a suitable action for each of the cells encountered due to drift. So, when the corridor becomes longer, more cells are encountered due to nondeterminism, and therefore the learning process takes more time to learn a suitable action for each of the encountered cells. Figure 5.3a suggests an almost linear relationship between corridor length and the empirical upper bound on the convergence-trial index based on the 0.9-quantile. In Section 5.2, we will see another example, where the relationship is not linear.

[22] Of course, it is impossible to simulate an infinite run in practice.

[23] In Appendix A, we give the precise definition of p-quantile that we have used in our analysis.

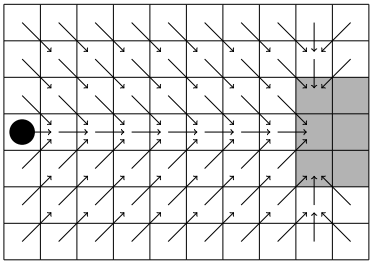

Figure 5.2: A corridor template for grid navigation tasks. The single start cell is marked with a black circle. The goal cells are shaded in gray. The parameterizable part is the horizontal distance between the start cell and the patch of goal cells. For completeness, the arrows represent a possible policy that illustrates reducibility of this kind of task. Although not shown, we recall that the finish action should be executed on the goal cells. The length of the depicted corridor is 7, which is the distance between the start cell and the patch of goal cells.

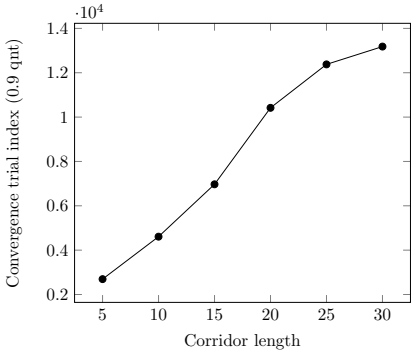

(a) Convergence-trial index related to corridor length.

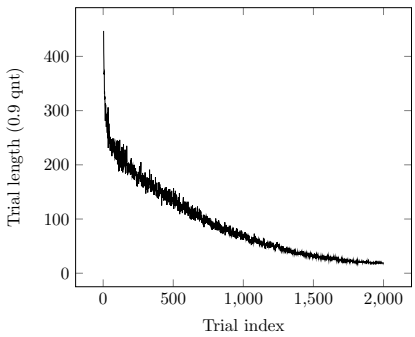

(b) Trial length related to trial index for the corridor task with length 10. We have excluded the first two trials to better show the shape of the curve.

Figure 5.3: Plots for the corridor grid navigation tasks.

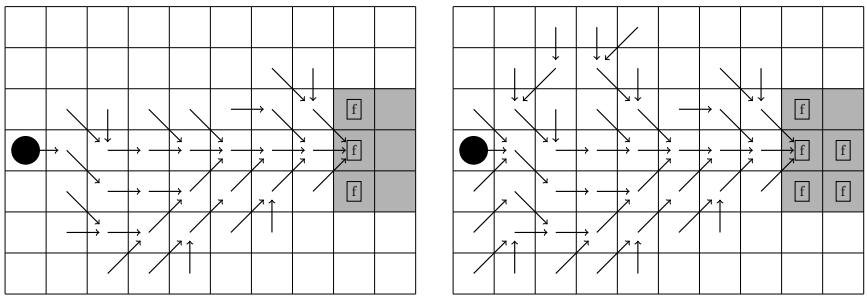

(a) Forward state set. (b) Backward state set.

Figure 5.4: Visualization of the forward state set and backward state set, from Section 4.2, of the final policy in one simulated run of the corridor navigation task shown in Figure 5.2. The states belonging to a set are marked with their action in the final policy; the empty grid cells are not part of the set. Each arrow shows only the main intended direction of the action; we do not show the left-rotating noise option and the right-rotating noise option. Goal cells that are assigned the finish action by the final policy are marked with the "f" sign.

Trial Length For a fixed grid navigation task, we also wanted to test if trial length decreases as the run progresses. A decreasing trial length would demonstrate that the learning algorithm is gradually refining the policy to go more directly to reward from the start states, avoiding cycles in the state space.

For the fixed corridor length of $l = 10$, we have simulated the first 2000 trials of 1000 runs; and, for each trial we have measured its length as the number of transitions. This gives a data matrix where each cell (i, j), with run index i and trial index j, contains the length of trial j in the simulated run i. For each trial index we have computed the 0.9-quantile of its length measurements in all runs. By plotting the resulting empirical upper bound on trial length against the trial index, we arrive at Figure 5.3b. We can see that trial length generally decreases as the run progresses. A similar experimental result is also reported by Potjans et al. (2011), in the context of neurons learning a grid navigation task.

Visualizing Forward and Backward We recall from Theorem 4.5 that, within the context of a specific task, the final policy π in a converging run satisfies the inclusion $\mathit{forward}(\pi) \subseteq \mathit{backward}(\pi)$. For the corridor in Figure 5.2, for one simulated run, we visualize the forward state set and the backward state set of the final policy in Figure 5.4. Note that the forward state set is indeed included in the backward state set. Interestingly, in this case, the simulated run has initially learned to perform the finish action in some goal cells, as witnessed by the backward state set, but eventually some of those goal cells are no longer reached from the start state, as witnessed by the forward state set.

5.2 Chain Tasks

In addition to the concrete grid tasks of Section 5.1, we have also considered a slightly more abstract form of task, that we call chain task. The general form

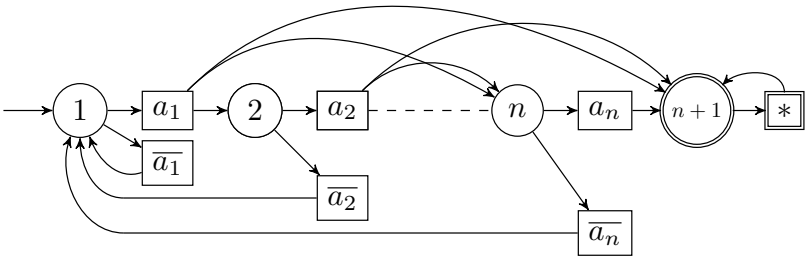

Figure 5.5: Template for chain tasks. Regarding notation, the symbol $*$ denotes the set of all actions $\{a_1, \ldots, a_n\}$, and the symbol $\overline{a_i}$ with $i \in \{1, \ldots, n\}$ stands for the set of all actions except action a_i. The graphical notation of tasks is explained in Figure 3.1.

of a chain task is shown in Figure 5.5. The parameter n tells us how long the chain is; the states of the chain are, in order, $1, \ldots, n$ and one final state $n+1$. To obtain reward from the start state, we should in the worst case perform all actions $a_1, \ldots, a_n$, in sequence, finished by one arbitrary action. For each $i \in \{1, \ldots, n\}$, the action a_i should be applied to state i. But there is forward nondeterminism that, for each pair (i, a_i) could send us to an arbitrary state later in the chain, closer to state $n+1$. Also, there are backward deterministic transitions that take us back to the start state whenever we apply the wrong action to a state. Formally, for a fixed value of $n \in \mathbb{N}_0$, we obtain a graph $T_n = (Q, Q_0, A, \mathit{rewards}, \delta)$, defined as follows: $Q = \{1, \ldots, n, n+1\}$, $Q_0 = \{1\}$, $A = \{a_1, \ldots, a_n\}$, $\mathit{rewards} = \{(n+1, a) \mid a \in A\}$, and regarding δ, we define

- for each $i \in Q$,

$$\delta(i, a_i) = \{i+1, \ldots, n+1\},$$
$$\delta(i, b) = \{1\} \text{ for each } b \in A \setminus \{a_i\}; \text{ and,}$$

- for the goal state $n+1$, we define $\delta(n+1, a) = \{n+1\}$ for each $a \in A$.

Note that for each $n \in \mathbb{N}_0$, the task T_n is reducible: conceptually, the reducibility iterations first assign action a_n to state n, then action a_{n-1} to state $n-1$, and so on, until state 1 is assigned action a_1.

On chain tasks, we have performed simulation experiments that are similar to the experiments on the grid navigation tasks in Section 5.1. These experiments are discussed next.

Convergence-trial Index Using the same experimental procedure as for grid tasks, we have simulated runs for some chain lengths. For each chain length separately, we have simulated 400 runs, and we have computed the 0.9-quantile of the measured convergence-trial indexes. By plotting the resulting empirical upper bound on the convergence-trial index against the chain length, we arrive at Figure 5.6a. We see that the convergence-trial index rises faster than linear in terms of the chain length; this is in contrast to Figure 5.3a for grid corridors. One possible explanation, is that the forward nondeterminism on the chain

causes each state to be visited less frequently in a simulated run. For a fixed length n, the effect is that a state $q \in \{2, \ldots, n\}$ could stay connected longer to a bad action, leading back to the start state 1. Of course, at the end of each trial, the start state should be connected to action a_1, because otherwise no progress could have been made; but the other states could in principle have any action. So, we might not yet encounter the final policy for a long time, as recognizable through the syntactic characterization of Theorem 4.5.

Trial Length Using the same experimental procedure as for grid tasks, for the fixed chain length of $n = 10$, we have simulated the first 2000 trials of 1000 runs, and we have computed the 0.9-quantile on trial length as explained for the grid corridor experiment. By plotting the resulting empirical upper bound on trial length against the trial index, we arrive at Figure 5.6b. Again, this figure suggests that the learning algorithm is able to gradually improve the policy over trials.

Form of the Final Policy We describe the form of the final policy, rather than visualizing it, because the form is very restricted. Consider the final policy π in a simulated run of the chain task with length $n \in \mathbb{N}_0$. In the simulation, we know that each trial should have ended with action a_1 assigned to state 1 because otherwise the state $n + 1$ could not have been reached. This property also applies to the convergence-trial, so $\pi(1) = a_1$. Hence, $\mathit{forward}(\pi)$ is the set of all states due to the forward nondeterminism along the chain. We recall from Theorem 4.5 that the final policy satisfies $\mathit{forward}(\pi) \subseteq \mathit{backward}(\pi)$. Therefore, $\mathit{backward}(\pi)$ is also the set of all states. We can now see that the final policy π has to satisfy $\pi(i) = a_i$ for all $i \in \{1, \ldots, n\}$. Towards a contradiction, let i be the largest state number for which $\pi(i) \neq a_i$. Then, using the iterations for computing $\mathit{backward}(\pi)$, we can see that $\{i + 1, \ldots, n + 1\} \subseteq \mathit{backward}(\pi)$ but $\{1, \ldots, i\} \cap \mathit{backward}(\pi) = \emptyset$.[24] This would be the desired contradiction, because $\mathit{backward}(\pi)$ is the set of all states.

6 Conclusion and Further Work

We have studied the fascinating idea of reinforcement learning in a non-numeric framework, where the focus lies on the interaction between the graph structure of the task and a learning algorithm. We have studied the graph property of reducibility, that implies the existence of a policy that makes steady progress towards reward despite nondeterminism in the task. Interestingly, reducibility, combined with a natural fairness assumption, enables our simple learning algorithm to learn the task. We have also characterized the final policy for converging runs, which allows the precise detection of the convergence-trial in simulations. We now discuss some avenues for further work.

Characterizing Learnable Tasks We have seen a sufficient property (Theorem 4.1) and necessary properties (Proposition 4.6) for tasks to be learnable under fairness. The gap between the sufficient and necessary properties seems strongly related to the unstable set of a task (see Section 4.3). Perhaps it is

[24] Here, the computation of $\mathit{backward}(\pi)$ is started with the set $\mathit{ground}(\pi) = \{n + 1\}$.

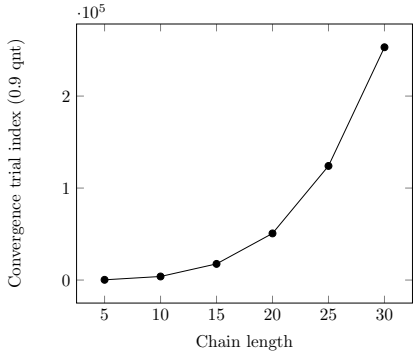

(a) Convergence-trial index related to chain length.

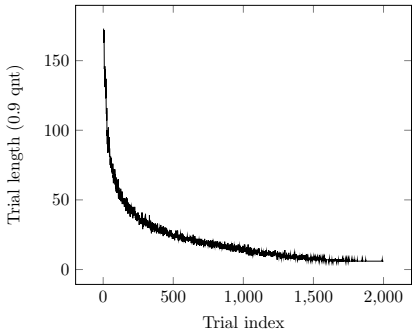

(b) Trial length related to trial index for the chain task with length $n = 10$. We have excluded the first two trials to better show the shape of the curve.

Figure 5.6: Plots for the chain tasks.

possible to characterize the tasks that are learnable under fairness by imposing some additional constraints on the manner by which the unstable set is connected to the reducible states, in addition to the already identified necessary properties.

Time Before Convergence Related to Remark 4.2 and to the simulations in Section 5, one could try to theoretically provide an upper bound on the convergence-trial index, for some class of tasks. We could also make assumptions regarding the probability distributions underlying the random actions proposed for a state and the choice of successor states when applying an action. For a given task, the result could be a probability distribution on the convergence-trial index, or on the total number of transitions before convergence.

Fading Eligibility Traces In some models of biologically plausible learning, the activation between any pair of connected neurons is represented by an eligibility trace (Sutton and Barto, 1998; Frémaux et al., 2013; Gerstner et al., 2014). When obtaining reward, the value of this trace is applied to the synaptic weight between the neurons. In realistic scenarios, the traces fade, with the advantage that a simulation or practical setup does not have to keep remembering all the information of past states before obtaining reward. The current article may be viewed as studying eligibility traces that are non-fading, because the working memory of the cycle-detection learning algorithm is effectively stored until the end of the trial. It appears interesting to let the working memory fade, perhaps by modeling the working memory as a first-in-first-out queue, where a newly entering state would remove the oldest state from the queue once a size limit has been reached.

Negative Feedback In this article, any path through the state space from a given start state to reward is good. However, in some applications we want to avoid certain states. For example, in a navigation task, an organism might want

to reach its nest from a certain starting location, but on the way to the nest the organism should avoid hazardous locations like pits or swamps. It appears interesting to formally investigate such cases. Information about hazards could be incorporated by extending the framework of this article with an explicit set *hazards* in tasks, which contains the state-action pairs which should be avoided; this is the opposite of the set *rewards*. The framework could further be extended to differentiate between trials that terminate with reward and trials that terminate with a hazard. A run could be said to convergence if eventually all trials terminate with reward, and all states eventually become stable.

Incomplete Information and Generalization The framework studied in this article provides complete information to the learning algorithm because each individual state can be mapped to its own action. In real-world applications, such as robot navigation (Thrun et al., 2005), the agent can only work with limited sensory information available in each time step. In that case, the agent should first build concepts for states in order to differentiate them. These concepts should be made by remembering sensory information over time, and in general multiple states will remain grouped together under the same concept because sensory information is not sufficiently accurate to differentiate between all states. This issue was also raised as an item for future work by Frémaux et al. (2013), who initially also have considered a framework in which complete information is available to the learning agent. There is ongoing work on partially observable tasks, see e.g. (Littman et al., 1995; Chatterjee and Chmelík, 2015).

Building concepts is related to the problem of generalization (Sutton and Barto, 1998) because in real-world tasks there might be too many states to store in the policy. It would be useful to collect states in conceptual groups and then assign an action to each group.

It appears interesting to formalize incomplete information and generalization in an extended version of the current framework, and to investigate sufficient and necessary properties for convergence.

References

Baader, F., Calvanese, D., McGuinness, D., Nardi, D., and Patel-Schneider, P., editors (2010). *The Description Logic Handbook, Second Edition: Theory, Implementation and Applications*. Cambridge University Press.

Baier, C. and Katoen, J. (2008). *Principles of Model Checking (Representation and Mind Series)*. The MIT Press.

Bonet, B. and Geffner, H. (2006). Learning depth-first search: A unified approach to heuristic search in deterministic and non-deterministic settings, and its application to MDPs. In *Proceedings of the Sixteenth International Conference on Automated Planning and Scheduling*, pages 142–151.

Brafman, R. and Tennenholtz, M. (2002). R-MAX - A general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3:213–231.

Chatterjee, K. and Chmelík, M. (2015). POMDPs under probabilistic semantics. *Artificial Intelligence*, 221:46 – 72.

Cormen, T., Leiserson, C., Rivest, R., and Stein, C. (2009). *Introduction to Algorithms, Third Edition*. The MIT Press.

Dayan, P. (1992). The convergence of TD(λ) for general λ. *Machine Learning*, 8(3–4):341–362.

Dayan, P. and Sejnowski, T. (1994). TD(λ) converges with probability 1. *Machine Learning*, 14(3):295–301.

Francez, N. (1986). *Fairness*. Springer-Verlag New York, Inc.

Frémaux, N., Sprekeler, H., and Gerstner, W. (2013). Reinforcement learning using a continuous time actor-critic framework with spiking neurons. *PLoS Computational Biology*, 9(4):e1003024.

Gerstner, W., Kistler, W., Naud, R., and Paninski, L. (2014). *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge University Press.

Geva-Sagiv, M., Las, L., Yovel, Y., and Ulanovsky, N. (2015). Spatial cognition in bats and rats: from sensory acquisition to multiscale maps and navigation. *Nature Reviews Neuroscience*, 16(4):94–108.

Hyndman, R. and Fan, Y. (1996). Sample quantiles in statistical packages. *The American Statistician*, 50(4):361–365.

Jaakkola, T., Jordan, M., and Singh, S. (1994). On the convergence of stochastic iterative dynamic programming algorithms. *Neural Computation*, 6(6).

Kearns, M. and Singh, S. (2002). Near-optimal reinforcement learning in polynomial time. *Machine Learning*, 49(2–3):209–232.

Littman, M., Cassandra, A., and Kaelbling, L. (1995). Learning policies for partially observable environments: Scaling up. In *Proceedings of the Twelfth International Conference on Machine Learning*, pages 362–370.

Maass, W. (2014). Noise as a resource for computation and learning in networks of spiking neurons. *Proceedings of the IEEE*, 102(5):860–880.

Markram, H., Gerstner, W., and Sjöström, P. (2011). A history of spike-timing-dependent plasticity. *Frontiers in Synaptic Neuroscience*, 3(4).

Potjans, W., Diesmann, M., and Morrison, A. (2011). An imperfect dopaminergic error signal can drive temporal-difference learning. *PLoS Computational Biology*, 7(5):e1001133.

Schultz, W. (2013). Updating dopamine reward signals. *Current Opinion in Neurobiology*, 23(2):229 – 238.

Schultz, W. (2015). Neuronal reward and decision signals: From theories to data. *Physiological Reviews*, 95(3):853–951.

Strehl, A., Li, L., and Littman, M. (2009). Reinforcement learning in finite MDPs: PAC analysis. *The Journal of Machine Learning Research*, 10:2413–2444.

Sutton, R. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1):9–44.

Sutton, R. and Barto, A. (1998). *Reinforcement Learning, An Introduction*. The MIT Press.

Thrun, S., Burgard, W., and Fox, D. (2005). *Probabilistic Robotics*. The MIT Press.

Tsitsiklis, J. (1994). Asynchronous stochastic approximation and Q-learning. *Machine Learning*, 16(3):185–202.

Vasilaki, E., Frémaux, N., Urbanczik, R., Senn, W., and Gerstner, W. (2009). Spike-based reinforcement learning in continuous state and action space: When policy gradient methods fail. *PLoS Computational Biology*, 5(12):e1000586.

Watkins, C. (1989). Learning from delayed rewards. *PhD thesis, Cambridge University*.

Watkins, C. and Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3–4):279–292.

Table A.1: Grid action offsets. For each movement action, we first give the main intended direction, followed by a left-rotating noise option and a right-rotating noise option.

Action	Possible offsets
finish	(0,0)
left	(-1,0), (-1,-1), (-1,1)
right	(1,0), (1,-1), (1,1)
up	(0,1), (-1,1), (1,1)
down	(0,-1), (1,-1), (-1,-1)
left-up	(-1,1), (-1,0), (0,1)
left-down	(-1,-1), (0,-1), (-1,0)
right-up	(1,1), (0,1), (1,0)
right-down	(1,-1), (1,0), (0,-1)

Appendix

A Examples

This Section contains additional details for the example tasks discussed in Section 5.

A.1 Grid Action Offsets

For our formalization of grid tasks, Table A.1 gives the offsets for each action.

A.2 Quantiles

We compute the quantiles with the statistics package R. For completeness, we recall here the definition of quantiles that we have used in our analysis; this definition is called Definition 1 by Hyndman and Fan (1996).

Let L be a nonempty list of numbers, possibly containing duplicates. Let $|L|$ denote the length of L. For an index $i \in \{1, \ldots, |L|\}$, we write $L[i]$ to denote the number at index i in L. We write $order(L)$ to denote the ordered version of L, where the numbers are sorted in ascending fashion; we keep duplicates, so $|order(L)| = |L|$.

Let $p \in \mathbb{R}$ with $0 < p < 1$. The p-quantile of a nonempty list L of numbers, denoted $Q(p, L)$ is defined as follows: denoting $j = \lfloor p \cdot n \rfloor$ where $n = |L|$,

$$Q(p, L) = \begin{cases} order(L)[j+1] & \text{if } p \cdot n - j > 0 \\ order(L)[j] & \text{if } p \cdot n - j = 0. \end{cases}$$

Intuitively, using the index $j = \lfloor p \cdot n \rfloor$ in the ordered list $order(L)$ is a good attempt at finding a number v such that a fraction p of L is smaller than or equal to v. The assumptions $0 < p$ and $p < 1$ ensure that we only apply valid indexes in the range $\{1, \ldots, |L|\}$ to the list $order(L)$.

A.3 Implementation Notes

The simulation was written with Java Development Kit 8. In our experimental results, we have measured only the number of transitions and the number trials in runs. Since no wall-clock time was needed, the exact running time of the simulation in seconds was not measured. During every transition, we have used a uniform sampling from the possible successor states. Concretely, given an array of successor states, we have used the function `Math.random()` to generate random indexes in this array, as follows: the double precision number returned by `Math.random()` can be converted into an integer by multiplying with the array length and subsequently truncating the resulting number with the `Math.floor()` function.